IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 21 July 2022, accepted 23 August 2022, date of publication 5 September 2022, date of current version 12 September 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3204027

RESEARCH ARTICLE

Fast Data Recovery for Improved Mobility Support in Multiradio Dual Connectivity

CARLOS PUPIALES[1,2], (Member, IEEE), DANIELA LASELVA[3], AND ILKER DEMIRKOL[4], (Senior Member, IEEE)

[1]Department of Network Engineering, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain
[2]FICA, Universidad Técnica del Norte, 100105 Ibarra, Ecuador
[3]Nokia, 9220 Aalborg, Denmark
[4]Department of Mining, Industrial and ICT Engineering, Universitat Politècnica de Catalunya, 08034 Barcelona, Spain

Corresponding author: Carlos Pupiales (carlos.pupiales@upc.edu)

This work was supported in part by the Secretaría de Educación Superior, Ciencia, Tecnología e Innovación (SENESCYT), Ecuador.

ABSTRACT Data aggregation is one of the crucial features of the 3GPP Multi-Radio Dual Connectivity (MR-DC) technology. However, mobility events and radio link failures, which may occur during the data aggregation, may pose challenges in meeting the latency, reliability, and throughput key performance indicators (KPIs). Unlike single connectivity, the user equipment (UE) in MR-DC operation can experience such events in either of the two base stations (BSs) serving the UE with MR-DC. In typical MR-DC deployments, these events occur more frequently in the BS acting as the secondary node (SN) since the SN operates at a higher frequency band. In this paper, we show that handovers (HOs) and signal blockage events that occur at the SN can create out-of-order data reception or losses at the UE's Packet Data Control Protocol (PDCP) layer, making the application stop receiving data for up to hundreds of milliseconds. Thus, challenging to meet the KPIs defined for such application. To mitigate this effect, we propose an intelligent and efficient mechanism that operates in the transmitting PDCP layer and significantly minimizes the data interruption periods suffered by the application when the UE aggregates data and HOs or failures of the SN occur. We use LTE/NR testbed experiments to show that the proposed mechanism achieves a high and stable aggregate throughput with near-zero interruption time and data reliability of at least 99.999% without transport layer retransmissions. The experiments are conducted for saturated TCP traffic and under link quality variations based on traces extracted from a Nokia-proprietary system-level simulator.

INDEX TERMS SN change, SCG failure, multi-connectivity, dual connectivity, 5G, data aggregation.

I. INTRODUCTION

Multi-Radio Dual Connectivity (MR-DC) [1] is a 3GPP solution that allows the UE to simultaneously communicate with two BSs in order to enhance the per-user data rate through data aggregation, improve the data reliability through data duplication, and/or provide mobility robustness through data offloading [2], [3], [4]. In MR-DC, the two BSs can use the same or different 3GPP radio access technologies (RATs) of Long Term Evolution (LTE) and 5G New Radio (NR). For ultra-reliable low-latency communications (URLLC) and enhanced mobile broadband (eMBB) use cases, MR-DC plays an important role in meeting certain KPIs defined for such use cases. For instance, MR-DC can help achieve a given reliability target without using time-consuming retransmission mechanisms such as hybrid automatic repeat request (HARQ) or automatic repeat request (ARQ) through data duplication. Alternatively, aggregating data from both BSs can improve the user data rate without significant hardware complexities.

In a typical MR-DC deployment, as illustrated in Fig. 1, one BS has macro cell coverage using frequencies in the range 1 (FR1), i.e., below 7.125 GHz, while the other BS has small cell coverage and may use frequencies in the range 2 (FR2), i.e., 24.25 GHz to 52.6 GHz, [4], [5], [6]. In such common scenario, user mobility causes the link using the FR2

The associate editor coordinating the review of this manuscript and approving it for publication was Tiankui Zhang.

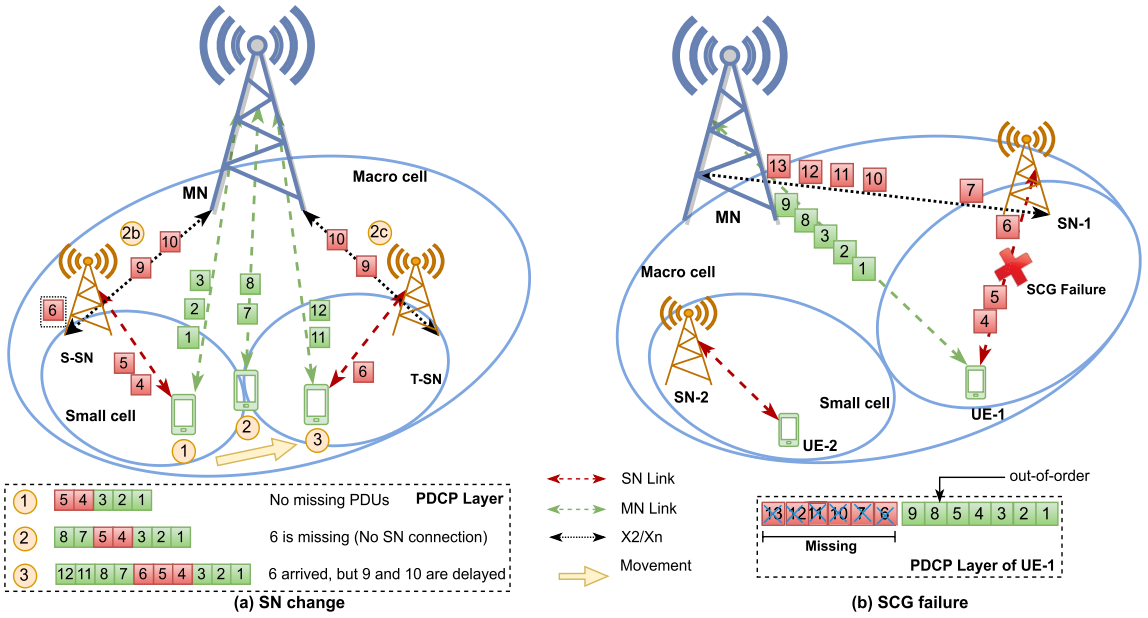

FIGURE 1. Out-of-order data reception caused by an SN change and SCG failure during data aggregation.

to suffer from signal blockages frequently, which might be caused either by loss in signal strength or because the UE transitions from the coverage area of one small cell to that of other small cell. In both cases, the UE losses connectivity with the corresponding BS. To recover the network connectivity, the UE or BS can perform a cell re-establishment in the former case, while the BS can trigger a HO for the latter case. Nevertheless, these procedures are time-consuming and cause the UE to stop receiving data from the small cell connection. The time the UE does not receive data is typically known as data interruption time, representing, from the physical layer perspective, a few dozens of milliseconds [7]. However, from the perspective of the upper layers, i.e., transport and application layers, the data interruption time may be much higher because of the delay added by the Medium Access Control (MAC), Radio Link Control (RLC), and PDCP layers [7], [8]. In this paper, we study the data interruption time from the perspective of the upper layers, which is more critical for the applications.

If the problems mentioned above occur while the UE aggregates data, the UE does not lose network connectivity completely. Indeed, it can still communicate via one of the BSs. Nonetheless, this can cause out-of-order PDCP packet data units (PDUs) arrivals, as illustrated in Fig. 1. In this scenario, the 3GPP-defined reordering mechanism at the UE's PDCP uses the PDU sequence number, a reordering timer, and a reordering window to wait for any delayed PDU to arrive and hence providing in-sequence data delivery to the upper layers [9]. Despite this effective mechanism, an inadequate configuration of the reordering timeout value makes the PDCP layer discard the delayed PDUs too early or buffer the PDUs for an excessively long amount of time [2], [3]. In both cases, they are probably triggering the upper layer retransmission mechanisms. For instance, unlike UDP, the TCP receiver requests data retransmission when it detects sequence gaps, a.k.a., fast retransmission, or when the already transmitted data has not been acknowledged during a given period at the TCP sender, a.k.a., retransmission by timeout. In these cases, the TCP receiver stops delivering new information to the application layer until the missing data is correctly received, increasing the data interruption time to several hundreds of milliseconds. On top of that, the aggregate throughput is seriously affected since TCP reduces its congestion window. In this situation, meeting the KPIs defined for reliability- and latency-constrained applications such as low latency eMBB is challenging [10]. Accordingly, in this paper, we show quantitatively the TCP performance degradation in such scenarios.

Despite the challenges that the data interruption time represent for the performance of MR-DC operation, the 3GPP has not defined any solution to tackle such a problem. Indeed, the SN change procedure [1], which is specified to manage the frequent changes of the SN, makes the UE stop communicating via the SN link until the change from the serving SN (S-SN) to the target SN (T-SN) is completed, as depicted in Fig. 1(a). Moreover, the typical way to recover the network connectivity from a radio link failure (RLF), i.e., the blockage of the radio link, in single connectivity (SC) operation is via a cell re-establishment procedure [11], [12]. However, when the SN link fails in MR-DC operation, a.k.a, secondary cell group (SCG) failure, such a solution is not specified by the 3GPP for MR-DC [8], [12], [13]. Therefore, the data buffered in the failing SN may be considered lost unless it can be transmitted to the UE using a new BS, as shown in Fig. 1(b).

Unfortunately, this requires a new data forwarding procedure, which the 3GPP has not considered. Note that this scenario can also occur if the SN change procedure fails.

Most of the relevant research studies in MR-DC such as [3], [14], [15], [16], [17], [18], [19], [20], [21], [22], [23], [24], [25] have proposed, on the one hand, methods to aggregate data considering that the UE is static and no failures occur. On the other hand, studies have proposed methods to provide mobility robustness. Yet, they do not consider the data aggregation or the data interruption time's effect on the application performance. For this reason, in this paper, we propose a *Fast Data Recovery method* (*FaRe)* that achieves near-zero data interruption time during SN change and SCG failure events. For this, the proposed *FaRe* method intelligently determines the PDCP PDUs split via the SN but not delivered to the UE and timely retransmit them via the master node (MN). Consequently, it leads to achieving high and stable aggregate throughput and lower latency compared to baseline strategies.

Using a mobile network testbed, we evaluate the performance of our *FaRe* method in a typical 3GPP-defined small cell scenario for saturated TCP traffic. Through testbed experiments, we show that when the *FaRe* is disabled at the MN, it can take approximately 74 ms and 24 ms for the first non-delivered PDU to arrive at the UE's PDCP layer when an SN change and SCG failure occur, respectively. On the other hand, with the proposed method, i.e., the *FaRe*, such a PDU is received in approximately 5-8 ms.

In summary, the contributions of this paper are manifold:

- We analyze the challenges that MR-DC faces to effectively aggregate data when the UE is moving and an SN change is required or an SCG failure occurs.
- We propose a fast data recovery method that allows the MN to timely retransmits the missing PDCP PDUs that may appear due to mobility events and/or radio link failures in MR-DC operation.
- We evaluate the proposed *FaRe* method using the LTE/5G-NR compliant Open Air Interface software and commodity hardware. The performance of the *FaRe* is evaluated against different benchmarking strategies using a 3GPP-defined small cell scenario in which the Channel Quality Indicator (CQI) resulting from mobility for the MN and SN are extracted from a Nokia-proprietary system-level simulator.

The rest of this paper is structured as follows. Section II introduces the main technical aspects and the challenges that a mobility scenario represents for the MR-DC operation. Moreover, in Section III, the problem addressed in this paper along with the most important related works are presented. In Section IV, the design principles and the functionality of the FaRe mechanism are described in detail. Section V describes the evaluation framework used for the validation of the FaRe and the results obtained by the evaluations of the FaRe and that of benchmarking strategies are shown in Section VI. The implementation impact of the proposed solution is detailed in Section VII. Lastly, Section VIII presents the conclusions of the work.

II. BACKGROUND ON SIGNALING MOBILITY MANAGEMENT

In the following, we present the most important aspects related to signaling mobility management for SC and MR-DC, considering downlink traffic.

A. MOBILITY WITH SC

In SC operation, the UE is connected to a single BS. Hence, the mobility events that may trigger a HO are handled by the BS the UE is connected to and the core network (CN). In the typical break-before-make HO used in LTE, the UE can experience, at the radio level, typical data interruption times of 15-50 ms, but delays of hundred of milliseconds can also occur [26]. Since the UE cannot communicate with the BS during the HO interruption time, some HO types have been proposed to solve this problem. For instance, the new RACH-less HO [27] can reduce the interruption time to 5 ms under certain circumstances [28] by avoiding the time-consuming Random Access (RA) procedure. In addition, the UE can use the Conditional HO (CHO) [27], [29] to trigger the HO if some conditions are met without involving the source cell, which is faster than the typical HO procedure. In case of HO failure, the UE usually performs a cell re-establishment, whose delay increases the HO interruption time. However, with the CHO, the UE can quickly apply the configuration of one of the already stored target BSs and synchronize with it. If the RA procedure is not successful, the UE can quickly start a new RA procedure with a new target BS, which reduces the HO interruption time in case of failures. Even though the new aforesaid HO types can help to reduce the interruption time, the signaling exchange required to support such HOs is still necessary between the UE and CN.

B. MOBILITY WITH MR-DC

In MR-DC operation, the UE maintains simultaneous connection with the MN and SN. Therefore, from the control plane (CP) perspective, the MN and SN independently handle the Radio Resources Management (RRM) procedures with the UE using their own Radio Resource Control (RRC) protocol. However, the connection to the CN is made via the MN only. On the other hand, the user plane (UP) traffic can be transferred to the UE via the MN, SN, or both, the decision of which mainly depends on the flow control algorithm used in the MN to split the incoming user's data, and the data radio bearer (DRB) configuration, i.e., split or direct DRB [1], [3].

Unlike SC, during mobility events affecting the SN in MR-DC, the user's traffic can be rapidly forwarded (a.k.a., fast switching) to a new BS, i.e., to the T-SN, if the UE is already synchronized with the BS. In this case, the HO interruption time to account depends only on the backhaul link, the delay of which can be considered as negligible from the UP perspective for some applications [21]. The Dual Active Protocol Stack (DAPS) [30] is a new HO solution to leveraging MR-DC and its dual protocol stack, in which the

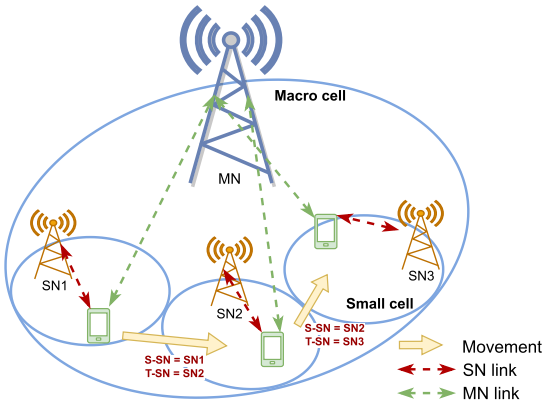

FIGURE 2. SN change events in a typical MR-DC deployment scenario.

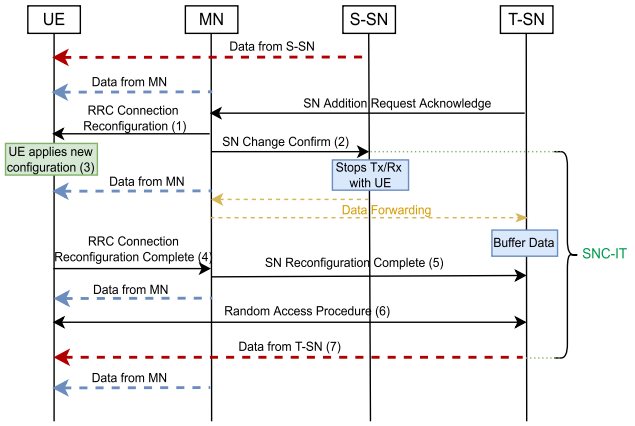

FIGURE 3. Signaling exchange for the SN change SN-initiated.

UE maintains the communication with the source BS until the UE has completed the RA procedure with the target BS. The UE releases the connection with the source BS after receiving an explicit notification from the target BS. Note that in scenarios where the UE is static, the DAPS HO is not generally useful. However, the MN can forward the data via the SN to provide load balancing instead of seamless connectivity in such cases.

According to the 3GPP [1], four mobility management scenarios are possible in MR-DC operation:

- SN change (MN/SN initiated). The MN is not modified and the UE continues with the MR-DC operation.
- Inter-Master Node handover with/without Secondary Node change. The MN is modified, but the UE can continue with the MR-DC operation if the SN is not changed.
- Master Node to eNB/gNB change. The UE switches to SC operation.
- eNB/gNB to Master Node change. The target MN adds an SN during the HO. Hence, the UE switches from SC to MR-DC operation during this procedure.

In a typical MR-DC scenario, as shown in Fig. 2, the HO events are more frequent for the SN than for the MN, especially in high-density areas. Hence, we focus on addressing the challenges of the first above-mentioned mobility scenario, i.e., the SN change, in this paper. However, our work can be extended to address the second mobility scenario as well, i.e., MN change.

During data aggregation, both BSs are simultaneously transferring UP traffic to the UE. Therefore, if the UE leaves the serving SN's coverage area, as depicted in Fig. 2, an SN change procedure [1], [12] is initiated either by the MN or SN to avoid losing network connectivity via the SN link. In this regard, when the new SN, i.e., the T-SN, confirms the allocation of radio resources for the UE, the MN sends either a *SN Release Request* or *SN Change Confirm* message to the S-SN indicating that it must stop the communication with the UE. Consequently, the UE can only communicate with a single BS, i.e., the MN, until the RA procedure with the T-SN is completed. Fig. 3 illustrates the signaling exchange required for the SN change procedure initiated by the SN. Considering the typical delays for the CP and UP of eMBB services [31], the backhaul link, and the different steps involved in the SN-initiated Change procedure, the communication between the UE and SN is interrupted for approximately 74 ms. Table 1 shows the typical contributors to the delay for this *SN Change Interruption Time (SNC-IT)*, initiated by the SN. For the MN-initiated case, the *SNC-IT* is slightly higher, i.e., 79 ms. Note that the *SNC-IT*, MN- or SN-initiated, can be longer if any of the steps cannot be completed, they take longer, or the backhaul latency is larger.

TABLE 1. Typical delays for SNC-IT, adapted from [26], [32].

Task	Typical Delay [ms]
(1) RRCConnectionReconfiguration	6
(2) SN Change Confirm	5
(3) UE applies new configuration	15
(4) RRCConnectionReconfigurationComplete	12
(5) SN Reconfiguration Complete	5
(6) Random Access Procedure	40
(7) Data from T-SN	1
SNC-IT = (1) + (3) + (4) + (5) + (6) + (7) - (2)	74

C. MCG/SCG FAILURES

Besides data interruptions experienced by the UE during SN change in MR-DC operation, interruptions due to RLFs in the Master Cell Group (MCG), i.e., MN, or SCG, i.e., SN, connections can also occur. In SC operation, when the UE declares an RLF, an RRC Connection Reestablishment procedure is initiated to recover the network connectivity [12]. However, in MR-DC operation, when the UE experiences an RLF on the MCG, the UE can trigger a fast MCG recovery procedure, if configured, or an RRC Connection Reestablishment procedure [1]. For the fast MCG recovery case, the UE stops the communication for all radio bearers configured in the MN and sends the *MCGFailureInformation* [12] message to the MN via the SN. Depending on the information

contained in the message, the MN can decide whether to change the MCG connection to a better cell or to release the MCG connection [1], [13].

For the case of an SCG failure caused by an RLF or an SN addition/change failure, the connection with the SN is suspended, and the UE notifies the MN with a *SCGFailure-Information* [12] message. Then, the MN decides whether to release, reconfigure, or change the failing SN. Note that since the MCG link is not affected, the UE can still receive the user's data via this link [13]. However, any pending packets at the SN may be lost during an SN change. Furthermore, when an MCG failure occurs and the fast MCG recovery is triggered, the UE will not receive the user's data via the MN link for approximately 30-70 ms [8], which is comparable with a typical HO interruption time. On the other hand, for an SCG failure, the data interruption can take significantly longer since the 3GPP does not specify any fast recovery method. It is worth mentioning that due to PDCP PDU losses or out-of-order data reception, the data interruption time experienced by the upper layers will be much higher than the one experienced by the physical layer [10].

III. PROBLEM DESCRIPTION AND RELATED WORK

During data aggregation in MR-DC, the out-of-order arrival of PDCP PDUs makes the receiver PDCP layer waits for any delayed PDU until the PDCP sequence number gap is filled or the reordering timer expires. During this period, the already received PDUs cannot be delivered to the upper layers, thus, the application stops receiving data. During SN-based HO or RLF, the out-of-order arrivals may be more frequent since the PDCP layer temporarily receives PDUs via only one BS. Having a large reordering timeout value can avoid discarding the delayed PDUs at the SN. Nevertheless, the application's data interruption time might also increase, making the application unable to achieve the required KPIs. For instance, low-latency eMBB applications such as Virtual Reality (VR) and Augmented Reality (AR) have a *motion-to-photon* latency requirement of less than 20 ms in order for the head's movement to match with the virtual scenery change [33]. In this regard, without considering the processing delays, their packet delay budget (PDB) is around 7-15 ms [33]. If during the transmission of this kind of data, a HO is required or a RLF occurs, the latency requirement of less than 15 ms is almost impossible to achieve with the current SN change procedure.

In a typical MR-DC deployment, i.e., FR1 + FR2 MR-DC, in which FR2 is used for capacity boosting, the UE will likely require numerous SN changes to satisfy the end-user reliability targets and provide seamless network connectivity via both BSs. In this regard, if an SN change is required when aggregating data via both BSs, the receiver PDCP layer will be impacted as it can receive PDUs from the T-SN only after the UE is successfully synchronized with the T-SN, as shown in Figs. 2 and 3. Note that it is not possible to use the DAPS HO in this case since both protocol stacks are already in use unless there is a third radio for MR-DC + DAPS. Additionally, when an SCG failure occurs due to an RLF or SN change failure, the UE will no longer receive PDUs via the SN link. In this case, the PDCP PDUs buffered in the SN's RLC buffer, i.e., the ones transmitted but not received at the UE, and the ones in-flight, i.e., the PDUs already split by the MN but which have not arrived at the SN yet, can be considered lost unless they are transmitted by a new BS.

Most of the available research efforts on MR-DC have focused on developing flow control solutions for data aggregation or methods to provide mobility robustness without data aggregation instead of reducing the application's data interruption experienced during HOs or signal blockages. For the former case, flow control mechanism mainly aim to maximize the user's aggregate throughput, reduce the end-to-end latency, maximize the throughput in one of the BSs, or achieve a minimum throughput for all users in both BSs [3], [14], [15], [16], [17], [18], [19], [20]. All theoretical and practical models presented in these works and their evaluations consider UEs without mobility and without service interruptions. Moreover, for the latter case, there have been studies that explore the use of MR-DC as an alternative to the legacy HO. In these studies, the authors state that MR-DC can reduce the HO failure probability, signaling exchange with the CN, the HO computational complexity, and HO completion delay [21], [22], [23], [24], [25]. Furthermore, these studies consider that before the HO, the UE already had CP connectivity via both BSs, i.e., the split bearer is configured in the MN and SN. However, the user's data is always transmitted via only one BS, i.e., the SN, upon triggering the HO. For this, the traffic is forwarded from the MN to the SN.

Furthermore, several studies show the capability of MR-DC to reduce the negative impact of link blockages on the performance of the user application's KPIs. For instance, in [34], the impact of various system parameters on the user's ergodic capacity for dense mmWave deployments is studied. Authors demonstrate that using multiple degrees of multi-connectivity, i.e., multiple radio link connections, helps to increase the achievable capacity by enabling backup connections. Additionally, in [35], authors indicate that having the UE with multi-connectivity reduces by up to seven times the denial of service and by up to ten times the dropping probability when static and dynamic blockages appear at the density of one blocker per square meter. Moreover, the theoretical framework presented in [36] suggests that under a high-density BSs deployment, extensive UE coverage, and short HO execution time, dual connectivity is sufficient to achieve the reliability target required in URLLC services in the presence of signal blockers. However, the multi-connectivity degree needed to support VR/AR services may be higher, especially in ultra-dense deployments. Authors in [37] state that blockages reduce the line of sight probability between the UE and BSs, implying that the UE has fewer available BSs to connect with in the area. This, in turn, increases the HO likelihood in BSs that use mmWaves. As a result, having multiple radio link

connections improves the link reliability under simultaneous mobility and blockage events.

Our literature survey shows the need for studying solutions to minimize the data interruption periods that the application suffers during data aggregation when SN change or SCG failure events occur. This problem has been addressed only from the radio level perspective, i.e., the physical layer, and in SC scenarios. Nevertheless, since the splitting and aggregation processes are done at the PDCP layer, an effective solution to the problem must include the perspectives of transport and application layers. Based on our knowledge, no work has studied the impact that mobility and signal blockages have on the performance of data aggregation with MR-DC, nor proposed solutions for such a problem. Actually, this topic has been included in the MR-DC enhancements plan for the future 3GPP Release 18 [10], [38], [39], [40], [41]. For this reason, we propose a fast data recovery method for MR-DC that intelligently and effectively minimizes the application's data interruption time as described in the following sections.

IV. FAST DATA RECOVERY ALGORITHM

A. TARGETED CHALLENGES

In MR-DC operation, data aggregation entails splitting the incoming user's traffic via the MN and SN and then reversely aggregating it at the receiver PDCP layer. A maximum aggregation benefit [42], i.e., the throughput gain obtained with MR-DC over SC, is achieved if the MN and SN maintain a continuous data flow with the UE [3]. For this, the RLC buffers of both BSs should have enough data to be transmitted at every transmission opportunity. In light of that, the flow control algorithm used in the MN should decide, based on each BS's radio link conditions and available radio resources, the number of PDCP PDUs to split via the MN and SN.

The different flow control algorithms available in the literature use up-to-date UE statistics, e.g., channel state information (CSI), from both BSs for the splitting decision making. However, during an SN change event, the UE is no longer connected to the corresponding SN; thus, there are no statistics to send back to the MN in order for the flow control mechanism to continue with the data splitting via the S-SN. A reasonable action, in this case, can be to stop the data splitting via the S-SN until the UE has completed the RA with the new BS, i.e., the T-SN. However, this can cause the T-SN to have no data to transmit once the synchronization is completed, reducing the aggregate throughput. On the other hand, maintaining the data splitting via the S-SN can increase the out-of-order arrivals or losses of PDUs, exacerbating the application's data interruption time, especially if the SN change fails. In this regard, whether to suspend or not the data splitting via the S-SN, at which stage of the SN change procedure to do it, and when to resume the splitting, if stopped, are significant aspects to consider during an SN change event.

Additionally, in an SCG failure event, the communication via the SN link is temporarily unavailable. Still, it can be resumed if the MN decides to keep or change the failing SN.

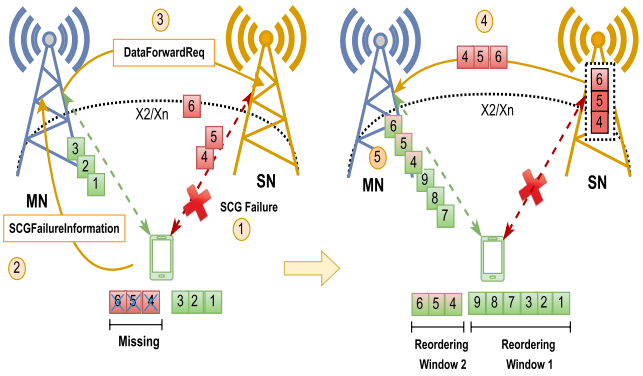

FIGURE 4. Data forwarding after an SCG failure event.

However, this procedure is time-consuming and can take hundreds of milliseconds, possibly making the PDCP reordering mechanism discard the PDCP PDUs pending at the SN since they will likely arrive at the UE within a different reordering window. In other words, the sequence number of the received PDU(s) will be lower than that of the last PDU delivered to the upper layers, as illustrated in Fig. 4. Moreover, since the non-delivered data is probably present in the failing SN's RLC buffer, it can be *re-routed* and transmitted by a new BS, e.g., via the MN. However, the data must be *re-routed* through a backhaul link with a non-zero delay, which makes it challenging to meet the latency requirements for some applications. For instance, according to [10], low-latency eMBB applications require that the data *re-routing* must be completed in less than 10 ms. In this regard, our proposed fast data recovery mechanism for MR-DC will address the abovementioned challenges.

B. FaRe DESIGN PRINCIPLES

Our *FaRe* mechanism aims to minimize the data interruption time that the application experiences during SN change or SCG failure events. To achieve that, the *FaRe* locally and temporarily stores the PDCP PDUs split via the SN. Therefore, the MN can timely retransmit the missing PDUs when one of the aforementioned events occurs. In this regard, the *FaRe* avoids the time-consuming higher layer data retransmissions, e.g., at the TCP level, which can be required when an SCG failure occurs. Likewise the *FaRe* avoids the slow data forwarding procedure required during an SN change. The *FaRe* works along with a flow control algorithm to facilitate the data splitting management, i.e., stop/pause/resume the splitting, during the SN change or failure events. The *FaRe* has three main functional stages: the buffering, the fast retransmission, and the splitting activation stages, which are depicted in Fig. 5 and described in the following.

1) BUFFERING STAGE

In MR-DC operation, the SN cannot independently manage the SN change or SCG failure events. Actually, the MN makes the SN release/change/reconfigure decisions. For example, if an SCG failure occurs, the failing SN cannot directly

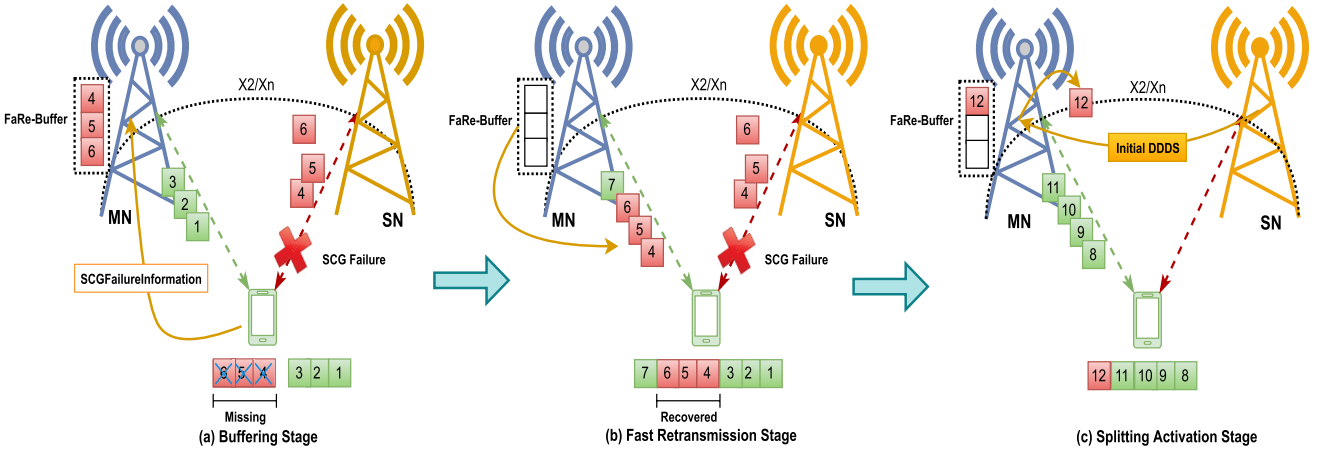

FIGURE 5. Functional stages of the *FaRe* mechanism.

forward the data present in its RLC buffer to a new BS. Instead, the MN has to request such a forwarding action, the data of which can be available for transmission in a new BS within $2 * backhaul\_delay$ ms, assuming equal backhaul delays between BSs and zero message processing delay. Hence, depending on the backhaul delay, achieving a given application's latency target can be challenging.

To tackle this challenge, the *FaRe* dispenses with the slow data forwarding procedure. Instead, it temporarily stores a copy of the PDCP PDUs split via the SN in a newly defined buffer, named *FaRe-Buffer*, which is placed at the MN's PDCP layer, as depicted in Fig. 5(a). These PDU copies are stored in a FIFO manner until the SN indicates that such PDUs have been successfully delivered, RLC AM case, or transmitted to the UE, RLC UM case. Note that the head of the buffer is the PDU copy with the smallest PDCP sequence number.

In addition, to the reduce the memory allocation requirements at the MN, the *FaRe* acknowledges the buffered PDUs using the highest transmitted/delivered PDCP sequence number information provided by the SN within the 3GPP-standardized Downlink Data Delivery status (DDDS) report [43]. Once the PDUs' transmissions have been acknowledged at the MN, the corresponding *FaRe-PDUs*, i.e., the PDCP PDU copies present in the *FaRe-Buffer*, are deleted. It is worth mentioning that the SN can periodically send the DDDS report through the *X2/Xn* interface, or the MN can explicitly request it via the same interface [43]. The Algorithm 1 describes the procedure performed by the MN to store and acknowledge the *FaRe-PDUs*.

Algorithm 1 FaRe Buffering Stage Algorithm

Input: PDCP PDU or DDDS report
Output: *FaRe-PDUs*, Acknowledged *FaRe-PDUs*

```
while FaRe is enabled do
    if PDCP PDU is forwarded to the SN then
        Make a copy of the PDU
        Place the copy in the FaRe-Buffer
    else
        Continue
    if DDDS Type 1 is received then
        ACK_SN = highest transmitted/delivered SN
        while FaRe-PDU ≤ ACK_SN do
            Delete the FaRe-PDU
    else
        Continue
```

2) FAST RETRANSMISSION STAGE

a: SN CHANGE SCENARIO

As we mentioned earlier, the MN should stop the data splitting via the SN during an SN change event to minimize the out-of-order arrivals, thus the application's data interruption time. In this regard, the *FaRe* communicates the flow control mechanism to stop the data splitting via the SN before the MN sends either the *SN Release Request* or *SN Change Confirm* messages. Actually, the MN can even stop the data splitting earlier if the radio link conditions experienced between the UE and S-SN are not favorable to maintain the connectivity.

Upon receiving the *SN Release Request* or *SN Change Confirm* messages, the S-SN stops communicating with the UE and releases the radio resources assigned to the corresponding UE [1]. For this reason, the S-SN prepares and sends through the *X2/Xn* interface a DDDS report that includes the latest delivered/transmitted PDCP PDUs and the indication that this report is the final one, i.e., the Final Frame Indication flag is activated. After receiving the final DDDS report, the MN acknowledges the corresponding *FaRe-PDUs* and updates the *FaRe-Buffer* accordingly. Afterward, all the remaining *FaRe-PDUs* are placed at the head of the RLC buffer to be transmitted in the next transmission opportunity. Consequently, these PDUs can rapidly be delivered to the UE using the MN's Uu instead of waiting for the completion of the SN change procedure. Note that if the RLC buffer contains a segmented Service Data Unit (SDU), the *FaRe-PDUs* are placed after this SDU.

Considering the current SN change procedure initiated by the SN as described in [1], a PDCP PDU that was ready to be transmitted by the S-SN but which was forwarded to the T-SN can arrive at the UE's PDCP in approximately 74 ms, as shown in Table 1. However, with the *FaRe* mechanism, the same PDU can arrive earlier. This delay is computed as

$$\Delta_{FaRe} = BH + UP_{delay} + \Delta_{proc} + n \times TTI, \quad (1)$$

where Δ_{FaRe} is the elapsed time from the S-SN stopping the communication with the UE until the first *FaRe-PDU* arriving at the UE's PDCP, BH is the backhaul delay, UP_{delay} is the user plane latency, e.g., 4 ms for eMBB [31], Δ_{proc} is the time that takes to process the DDDS report and place the *FaRe-PDUs* in the RLC buffer, n is the number of elapsed TTIs until the first *FaRe-PDU* is scheduled for transmission, and TTI is the Transmission Time Interval, e.g., 1 ms for eMBB. Therefore, assuming $BH = 5$ ms, $UP_{delay} = 4$ ms, Δ_{proc} as negligible, $TTI = 1$ ms, and $n = 1$, i.e., the PDU is scheduled in the next TTI, the Δ_{FaRe} is 10 ms. This delay is significantly lower than the one achieved with the current SN change procedure, i.e., 74 ms.

b: SCG FAILURE CASE

Upon receiving the *SCGFailureInformation* [12] message from the UE, the *FaRe* communicates the flow control mechanism to stop the data splitting via the SN. Afterward, based on the PDCP Status Report [9] included by the UE in the *SCGFailureInformation*, the PDUs that have not been received at the UE are placed at the head of the RLC buffer so that they can be transmitted in the next transmission opportunity. The novel inclusion of the PDCP Status Report within the *SCGFailureInformation* message helps the MN to recover those PDUs that were erroneously received at the UE, which is common during poor radio link conditions like those experienced before the UE declares an RLF. The elapsed time from the MN receiving the *SCGFailureInformation* message until the UE's PDCP receiving the first *FaRe-PDU* is denoted as Δ^*_{FaRe} and computed as

$$\Delta^*_{FaRe} = UP_{delay} + \Delta_{proc} + n \times TTI. \quad (2)$$

Assuming $UP_{delay} = 4$ ms, $FaRe_{proc}$ as negligible, $TTI = 1$ ms, and $n = 1$, the Δ^*_{FaRe} is approximately 5 ms.

It is worth mentioning that there is no recovery mechanism in the literature that deal with the SCG failure problem. However, a possible approach could be to request data forwarding from the failing SN. In this case, the time that may take for the first non-delivered PDU to arrive at the UE's PDCP is denoted as DF and computed as

$$DF = 2 \times BH + UP_{delay} + RLC_{delay} + DF_{proc}, \quad (3)$$

where RLC_{delay} is the RLC buffering delay experienced at the moment that the forwarded PDUs ingress to the RLC buffer and DF_{proc} is the processing delay. Assuming $BH = 5$ ms, $UP_{delay} = 4$ ms, DF_{proc} as negligible, and an RLC_{delay} of 10 ms, the DF is approximately 24 ms. Note that the RLC_{delay} will vary depending on the flow control algorithm in use. For instance, the CCW flow control algorithm limits the RLC buffering delay to 20 ms [3]. Hence, assuming an RLC_{delay} at the 50% of its limit, i.e., 10 ms, is a conservative approximation to the reality.

The *Fast Retransmission Stage* functionality of the *FaRe* is presented in Algorithm 2 and is described in the following.

- The MN stops splitting PDCP PDUs via the BS acting as SN when it receives either the *SN Addition Request Acknowledge* or *SCGFailureInformation* message (Lines 2-3 and 14-15, respectively).
- After receiving the *SN Addition Request Acknowledge* message, the MN waits for the DDDS report with *FinalFrameInd* = 1. Meanwhile, it acknowledges the *FaRe-PDUs* and updates the *FaRe-Buffer* with every received DDDS report (Lines 4-8).
- Once the DDDS report with *FinalFrameInd* = 1 is received, the MN places at the head of its RLC buffer all the *FaRe-PDUs* that were not acknowledged. Subsequently, the MN flushes the *FaRe-Buffer* (Lines 9-13).
- After receiving the *SCGFailureInformation* message, the MN reads the PDCP-related information, acknowledges the *FaRe-PDUs*, and updates the *FaRe-Buffer*. If such a information includes the non-received PDUs, the corresponding *FaRe-PDUs* are not deleted from the *FaRe-Buffer*. Afterward, the *FaRe-PDUs* that were not acknowledged are placed at the head of the RLC buffer. Finally, the MN flushes the *FaRe-Buffer* (Lines 17-28).
- The MN resumes the data splitting via the SN once it receives an initial DDDS report from the T-SN or a new BS acting as SN (Lines 26-27).

3) SPLITTING ACTIVATION STAGE

As indicated in Section IV-B2, the *FaRe* makes the flow control to stop forwarding PDCP PDUs via the SN to avoid out-of-order deliveries that increases the application's data interruption time. To avoid the sharp drop in the aggregate throughput during the SN change, the T-SN should start transmitting the user data to the UE as soon as the RA procedure is completed. Nevertheless, having the T-SN's RLC buffer with sufficient data to be transmitted in the next transmission opportunity without increasing the data interruption time is challenging. For instance, if the SN change procedure fails, the PDCP PDUs present in the T-SN's RLC buffer should be forwarded to a new T-SN. Nevertheless, this time-consuming procedure may cause the aforesaid PDUs to arrive at the UE after the PDCP reordering timer expires. Hence, as they no longer belong to the reordering window provided to the upper layers, they are discarded [9], increasing the application's data interruption time, especially for TCP-based applications.

Since the *FaRe* timely transmits the non-delivered PDUs via the MN's Uu and makes the flow control to operate in SC-like mode during the SN change event, the receiving PDCP layer never stops sending data to the upper layers. Therefore, the throughput never goes down to zero. For this reason, the MN only resumes the data splitting via the SN

Algorithm 2 FaRe: Fast Retransmission Stage

Input: *SN Addition Request Acknowledge or SCGFailureInformation* messages
Output: *FaRe-PDUs* delivered to the UE's PDCP

```
while FaRe is enabled do
    if SN Addition Request Acknowledge received then
        Notifies the flow control to stop data splitting via SN
        while DDDS with FinalFrameInd = 0 do
            ACK_SN = highest transmitted/delivered PDU
            while FaRe-PDU ≤ ACK_SN do
                Acknowledge FaRe-PDUs
                Update the FaRe-Buffer
        if Head's RLC SDU is segmented then
            Place the FaRe-PDUs after the segmented RLC SDU
        else
            Place the FaRe-PDUs before the RLC SDUs
        Flush the FaRe-Buffer
    else if SCGFailureInformation received then
        Notifies the flow control to stop data splitting via SN
        Read the PDCP Status Report
        ACK_SN = First Missing PDU
        while FaRe-PDU < ACK_SN do
            Acknowledge FaRe-PDUs except the non-received PDUs included in the bitmap
            Update the FaRe-Buffer
        if Head's RLC SDU is segmented then
            Place the FaRe-PDUs after the segmented RLC SDU
        else
            Place the FaRe-PDUs before the RLC SDUs
        Flush the FaRe-Buffer
    else if Initial DDDS received then
        Notifies the flow control to resume data splitting via SN
    else
        Keep acknowledging FaRe-PDUs
```

when it receives from the T-SN the initial DDDS report [43], which indicates that the UE has completed the RA procedure and it is ready to receive new data.

Due to the T-SN's RLC buffer is empty, it cannot send to the MN any flow control statistic in order for the flow control algorithm to compute the number of PDCP PDUs to split via the SN. To overcome this challenge, the *FaRe* indicates the initial values for the variables that the flow control algorithm uses for its splitting decisions. For instance, the CCW flow control algorithm uses the MAC SDU sizes and RLC buffering delay statistics for its data splitting decisions [3]. Similarly, the Delay-based flow control algorithm uses the PDCP PDU transmission delay computed from the UE status report [16]. In this regard, considering these two state-of-the-art flow control algorithms, the above mentioned initial values are given by

a: THE CCW CASE

$$SDU_{DC}^{initial} = \min(SDU_{DC}^{1}, \ldots, SDU_{DC}^{k}), \quad (4)$$

$$RLC_{delay} = 0, \quad (5)$$

where $SDU_{DC}^{initial}$ is the initial value to use for the MAC SDU size variable, SDU_{DC}^{1} is the first MAC SDU size statistics received from the S-SN right after the MN sends the *SN Change Required* or *SN Addition Request* message, SDU_{DC}^{k} is the last MAC SDU size statistics received from the S-SN just before the MN sends the *SN Change Confirm* or *SN Release Request* message, and RLC_{delay} is the initial value to use for the RLC buffering delay. It is worth mentioning that during a HO, the quality of the radio link conditions does not allow the UE to achieve high data rates. However, in MR-DC operation, the BSs must assure a minimum of radio resources for the UE to achieve a given minimum data rate. For this reason, the *FaRe* uses the minimum MAC SDU size statistics received during the period mentioned above as the initial value for the MAC SDU size variable. Likewise, since the T-SN's RLC buffer is empty after the RA procedure, the *FaRe* indicates the CCW to use 0 ms as the initial value for the RLC buffering delay variable.

b: THE DELAY-BASED CASE

$$PDU_{delay} = 0, \quad (6)$$

where PDU_{delay} is the PDCP PDU transmission delay. Since during and right after the completion of the RA procedure, the T-SN has no data in its RLC buffer, the *FaRe* indicates the Delay-based to use 0 ms as the initial value for the PDCP PDU transmission delay.

Regardless of the flow control algorithm used by the MN to split the incoming data, the *FaRe* indicates the flow control algorithm to use the initial values for their variables until the MN receives up-to-date statistics from the T-SN. Likewise, the *FaRe's Buffering Stage* is initiated as soon as the incoming traffic is split via the SN link. Note that after an SCG failure event, the UE switches to SC operation. However, the UE may recover the MR-DC operation, thus the data aggregation, if the MN initiates the procedure to re-establish the connection with the failing SN [12]. On the contrary, if the MN decides to add a new SN, the data aggregation is started from scratch. Thus, the *Traffic Activation Stage* is not applicable in this case. The splitting activation stage is presented in the Algorithm 3.

V. EVALUATION FRAMEWORK

To validate our proposed *FaRe* mechanism, we implemented a Dual Connectivity (DC) [27] solution on a LTE/NR testbed developed using the Open Air Interface Software (OAI) [44]. The testbed is based on the split DRB architecture and implements the user plane functionalities of DC detailed in [9].

Algorithm 3 FaRe: Splitting Activation Stage

Input: *SN Addition Request or SN Change Required* messages

Output: Initial values for the flow control variables

while FaRe enabled **do**
  if *SN Addition Request* **or** *SN Change Required* sent **then**
    if CCW enabled **then**
      Set $SDU_{DC}^{initial} = 0$
      Set $RLC_{delay} = 0$
      while *SN Release Request* **or** *SN Change Confirm* has not been sent **do**
        if Flow control statistics received **then**
          if $SDU_{DC}^{initial} \leq SDU_{DC}$ **then**
            $SDU_{DC}^{initial} = SDU_{DC}$
          else
            Maintain the old $SDU_{DC}^{initial}$
    else if Delay-based enabled **then**
      Set $PDU_{delay} = 0$
    else
      Set the flow control variables accordingly.
  if *initial DDDS* report received **then**
    if CCW enabled **then**
      Set $SDU_{DC} = SDU_{DC}^{initial}$
      Set $D_q = 0$
    else if Delay-based enabled **then**
      Set $PDU_{delay} = 0$
    else
      Set the flow control variables accordingly.
    Indicates the flow control to resume the splitting via the SN

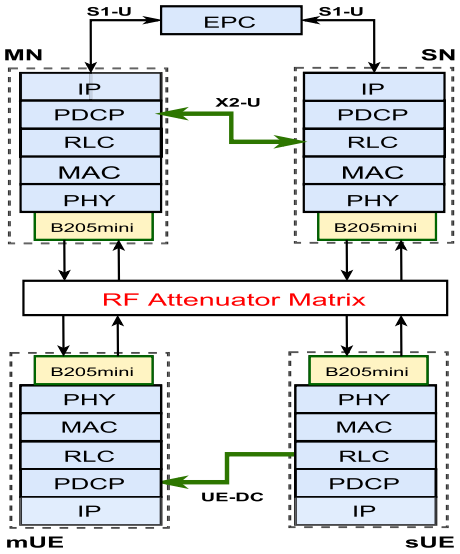

FIGURE 6. Testbed architecture [3], [45].

TABLE 2. General configuration for the BSs.

Parameter	Value
Duplex Mode	FDD
E-UTRA Band	7
DL Frequency for MN	2.68 GHz
DL Frequency for SN	2.63 GHz
Bandwidth for MN	5 MHz
Bandwidth for SN	10 MHz
RLC mode	UM

Further details of the testbed implementation can be found as part of our previous works in [3] and [45]. The testbed setup and the evaluation scenarios are given in the following.

A. TESTBED SETUP

Our DC testbed setup is created on different hosts of the ORBIT testbed [46], as illustrated in Fig. 6. Since the OAI's UE software does not implement the protocol stack required to support DC operation, i.e., one common PDCP layer and two independent RLC, MAC, and PHY layers, we use two OAI UE instances, i.e., mUE and sUE, to represent the UE in DC operation. In this regard, the hosts that represent the UEs, BSs, and CN are connected using a Gigabit-Ethernet switch. Additionally, every BS and UE host uses a Software-Defined Radio (SDR), of the model Ettus USRP B205mini, connected through a programmable JFW 50PMA-012 RF attenuator matrix [46]. Note that each pair of BS-UE has an isolated RF path since the SDRs are electromagnetically isolated from external signals. Moreover, the OAI software for the CN, BSs, and UEs runs over Ubuntu 16.04.1 with 4.15.0-52-low-latency kernel. The hosts have an Intel Core i7-4770 CPU @ 3.4GHz processor and 16 GB of RAM. Further configurations and parameters used in this study for the BSs are illustrated in Table 2.

B. BENCHMARKING STRATEGIES

As indicated in Section III, the problem of the data interruption at the application level caused by SN-based HO/RLF while the UE aggregates data has not been addressed yet in the available literature. Therefore, to validate the performance of the proposed *FaRe* mechanism, we have implemented different possible strategies to follow in case an SN change is required or an SCG failure occurs. For both cases, the baseline strategies use the state-of-the-art CCW flow control algorithm [3]. The CCW dynamically splits the incoming UP traffic via the MN and SN according to the average MAC SDU sizes and RLC buffering delays experienced in both BSs. To maximize the average aggregate throughput and achieve a stable value, the CCW maintains the RLC buffers of both BSs with enough data to prevent link starvation while avoiding congestion [3].

1) SN CHANGE SCENARIO

In the first strategy, *Baseline_1*, the MN, just before sending the *SN Change Confirm or SN Addition Request* messages,

resets the flow control variables related to the MN and S-SN, and makes the flow control operate under ''initial conditions'' until up-to-date statistics are received from both BSs. Operating under ''initial conditions'' implies for the CCW to split the data in a Round-Robin manner until the RLC buffering delay and MAC SDU size statistics from both BSs are known [3]. Note that the MN never stops splitting the incoming UP traffic via both BSs.

For the second strategy, *Baseline_2*, the MN stops the data splitting via the S-SN before sending the *SN Change Confirm or SN Addition Request* messages. The splitting is resumed when the MN receives the ''initial'' DDDS report from the T-SN, which indicates that the UE is successfully synchronized. Note that those PDUs that were not transmitted by the S-SN are forwarded to the T-SN as defined by 3GPP in [1]. Since the MN resumes the splitting only if the UE is synchronized with the T-SN, possible data losses or out-of-order deliveries caused by RA or RRC reconfiguration failures are avoided.

In the third strategy, *Baseline_3*, the MN resets the flow control variables related to the S-SN and suspends the data splitting via the S-SN before sending the *SN Change Confirm or SN Addition Request* messages. The MN resumes the data splitting when it receives up-to-date statistics from the T-SN. If the S-SN has no data to forward to the T-SN, the latter has no data to transmit to the UE once the RA procedure is completed. Thus, the MN will not resume the data splitting.

2) SCG FAILURE SCENARIO

For the baseline strategy, *Forward_Req*, upon receiving the *SCGFailureInformation* message, the MN requests the failing SN to forward the non-transmitted/delivered PDCP PDUs so that they can be transmitted via the MN's Uu. When the requested PDUs arrive at the MN, they are transmitted after the RLC SDUs present in the RLC buffer. Moreover, since during a data session the UE may experience SN change events before the SCG failure event occur, the *Forward_Req* uses the *Baseline_2* strategy to manage such HO events. This corresponds to a default implementation that follows the 3GPP specifications [1], [12].

C. PERFORMANCE METRICS

To evaluate the performance of our *FaRe* mechanism on reducing the application's data interruption time during SN change and SCG failure events, we use the following metrics:

1) AGGREGATE THROUGHPUT

We evaluate the aggregate throughput obtained at the transport layer with every Baseline strategy and the *FaRe* after a data session of 30 seconds for the SN change case, and 20 seconds for the SCG failure case. In this regard, we evaluate how the SN changes and SCG failure events impact on the instantaneous and average aggregate throughputs. Additionally, for a better illustration of the short-term impact of aforementioned events, the obtained aggregate throughput samples are collected every 100 ms.

2) THROUGHPUT VARIANCE

To evince and compare the variability of the throughput caused by SN change and SCG failure events in the entire data session, we use the Variance Ratio (R_{var}) [21], which is defined as

$$R_{var} = \frac{\delta_{T_{DC}}}{T_{DC}}, \tag{7}$$

where T_{DC} is the average aggregate throughput obtained by the application at the end of an experiment and $\delta_{T_{DC}}$ is the standard deviation of T_{DC}. Note that high values of R_{var} indicate significant throughput instability, such as long periods of zero throughput or short periods with very high throughput peaks.

3) DATA RELIABILITY

When a UE aggregates data, the main goal is to maximize the obtained throughput. However, achieving a given reliability target while maximizing the throughput may be challenging for some applications during SN change or SCG failure events. In this regard, we evaluate the reliability obtained at the PDCP level with and without the use of the *FaRe* mechanism. For this, we compare the number of PDCP PDUs that leave the MN's PDCP layer with the PDCP PDUs received in the mirroring layer at the UE during the entire data session. The PDCP reliability (R_{PDCP}) is defined as

$$R_{PDCP} = \frac{PDUs^{RX}}{PDUs^{TX}} \times 100\%, \tag{8}$$

where $PDUs^{RX}$ is the number of PDCP PDUs that are successfully received at the UE, and $PDUs^{TX}$ is the number of PDCP PDUs that are split by the MN and leave the PDCP layer to be transmitted via either BS.

4) DATA INTERRUPTION TIME

When the UE aggregates data, the interruption time experienced at the transport and/or application layers is influenced by the out-of-order arrival of PDCP PDUs or PDCP PDU losses. In this regard, the interruption time increases while the PDUs spend more time in the PDCP reordering buffer. Unlike UDP, TCP is a reliability-oriented protocol, so it must provide in-sequence delivery to the application. Therefore, if TCP sequence gaps are detected, the application will not receive data until the lost packet is correctly recovered by TCP. For this reason, we measure the elapsed time the transport layer stops receiving data during the data session. Note that *iperf3* measures the throughput at the transport layer, but these results also represent the data interruption time at the application level.

D. EVALUATION SCENARIO

To recreate a mobility scenario on the DC testbed described in Section V-A, we use the signal-to-interference-plus-noise-ratio (SINR), channel quality indicator (CQI), and reference signal received power (RSRP) traces extracted from the Nokia-proprietary system-level simulator for the MN and SN

using a 3GPP-defined DC scenario detailed in [47]. Note that the system-level simulator is entirely in alignment with 3GPP simulations and modeling guidelines. In addition, the SINR values from the aforementioned traces are measured every 5 ms. In this regard, each time a CQI is to be reported by the OAI UE, the next trace value is used. Hence, each OAI BS independently receives from the OAI UE, i,e., mUE and sUE stacks in our setup, the CQI value using the CSI report every 5 ms.

The 3GPP-defined DC scenario states nine small cells with inter-site distance (ISD) of 58 m, which are uniformly distributed across the macro cell area with no coverage gaps. To evaluate intra-MN MR-DC HO scenarios, the UE never leaves the macro cell coverage area in our evaluations, i.e., the UE is always connected to the same MN. In such a scenario, the UE moves at 30 Km/h following a single trajectory across the different small cells deployed in the macro cell coverage area. At the same time, the network is loaded with stationary UEs to create interference to have a realistic SINR. Furthermore, a 3 dB offset is considered for the HO decision, i.e., the A3 event, thus avoiding ping pongs between neighbor cells. Under the above-indicated conditions, during a data session of 30 seconds, the UE experiences nine SN change events.

In a typical MR-DC deployment, higher data rates in the small cell are achieved than in the macro cell thanks to higher bandwidth values assigned to the small cells. Therefore, we define a scenario where the MN and SN use different channel bandwidths, i.e., 5 MHz and 10 MHz, respectively. We only use this combination of bandwidth values because the OAI UE does not achieve a stable throughput with other values, e.g., 20 MHz [48]. In addition, since in our testbed, a single UE is connected to both BSs, all the physical resource blocks (PRBs) available for each channel bandwidth are assigned to the UE, i.e., 25 PRBs for the MN and 50 PRBs for the SN. In this scenario, we evaluate the performance of user traffic that uses the TCP as transport layer protocol during SN change and SCG failure events using the metrics described in Section V-C and the benchmarking strategies described in Section V-B.

Additionally, since the data interruption time may increase with the value configured for the PDCP reordering timer (t-Reordering), the TCP traffic is also evaluated using different reordering timeout values. For the evaluations, the downlink TCP traffic is generated using the *iperf3* tool, in which the *server* runs at the mUE's host, where the UE's PDCP layer resides, and the *client* at the EPC's host. Note that the *iperf3* intents to saturate the link to measure the maximum achievable throughput. Table 3 summarizes the scenario setup and provides additional configuration details.

TABLE 3. Scenario evaluated summary.

Parameter	Value
Single experiment duration	30 sec for SN change, 20 sec for SCG failure
Channel Bandwidth	$MN_{BW} = 5$ MHz $SN_{BW} = 10$ MHz
PDCP t-Reordering	100, 150, 200, 300 ms
PDCP reordering window	2048 PDUs
One-way backhaul latency	5 ms [26], [49]
Traffic Type	Saturated TCP
MR-DC architecture option	LTE-DC
3GPP-defined DC scenario	9 small cells per macro cell coverage area
CCW's configuration	$T_{CCW} = 5$ms, $D_q^* max = 20$ ms, and $\alpha = 0.3$ [3]

VI. RESULTS AND DISCUSSION

This section shows and discuses in detail the results obtained using the framework presented in Section V. According to the OAI specifications, the downlink throughput achieved for LTE in SC operation under an ideal scenario, i.e., highest CQI value and no packet losses, is 16-17 Mbps with a channel of 5 MHz and 34-35 Mbps with a channel of 10 MHz [48]. However, as we mentioned in Section V-D, the CQI values collected on the traces for the UE-MN and UE-SN connections vary depending on the UE's location concerning the BS in the macro and small cells.

A. AGGREGATE THROUGHPUT

Contrary to the throughput obtained in a static and/or ideal scenario, i.e., no losses and highest CQI, the aggregate throughput obtained in a mobility scenario is affected by the number of HO events, radio link conditions' variability, and the approach used to minimize the application's data interruption time. To illustrate the variability of the aggregate throughput, we present the results using the cumulative distribution function (CDF) obtained for 20 different experiments. The throughput samples for each experiment are collected every 100 ms for a total duration of 30 seconds in the SN change scenario and 20 seconds in the SCG failure scenario. The results presented in Figs. 7 and 9 do not include the throughput values obtained during the TCP' slow start stage, i.e., the first 3.5 seconds of the experiment duration in our case, to avoid confusing them with those obtained during SN change and SCG failure events. It is worth mentioning that random events such as HARQ retransmissions or packet losses may appear in our LTE/NR testbed, creating differences in the throughput values obtained across the experiments.

1) SN CHANGE SCENARIO

In this subsection, we evaluate the performance of the *FaRe* mechanism against the Baseline strategies described in Section V-B. Fig. 7 illustrates the CDF of the aggregate throughputs obtained using two *t-Reordering* values, i.e., 100 and 300 ms, which help us to visualize the behaviour of the aggregate throughput based on two different configurations.

The *Baseline_1* strategy achieves the worse performance among the compared methods. In this case, Fig. 7 shows that the probability of having zero throughput values is more significant than in the other cases. Indeed, the probability is around 1% only when the *t-Reordering* is 100 ms. In the

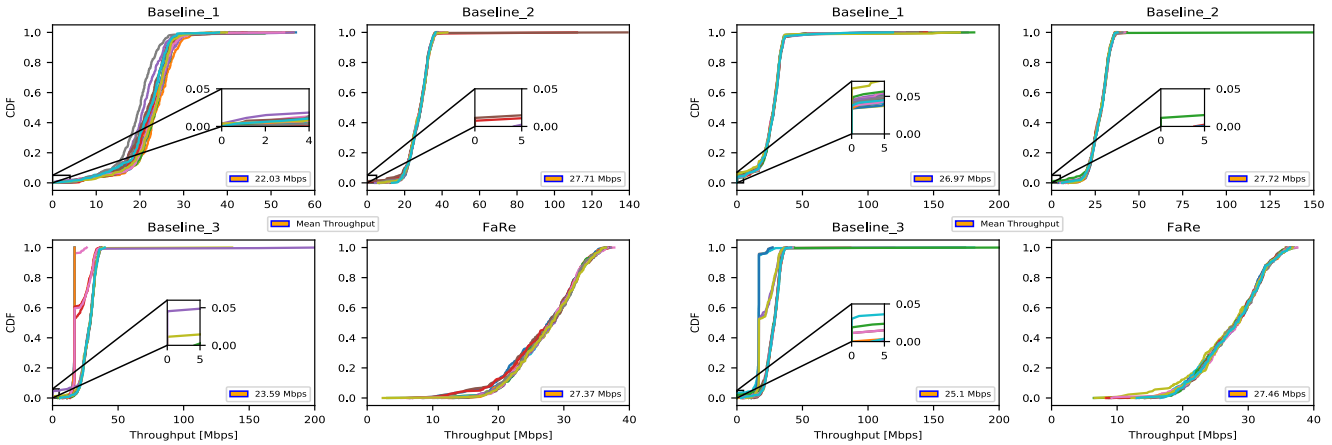

(a) CDF throughput for t-Reordering = 100 ms

(b) CDF throughput for t-Reordering = 300 ms

FIGURE 7. CDF of the aggregate throughputs obtained using different t-Reordering values for the SN change scenario.

other cases, the probability goes up to 4-6%. As a consequence of the periods with zero throughputs, the *Baseline_1* reaches high peak values of approximately 180 Mbps, which in turn harms the application's performance. In addition, even though with the other Baseline strategies, the probability of achieving zero throughput values drops to 1-2%, they still are not capable of achieving a stable throughput. Actually, abnormal peak values of more than 150 Mbps are common with the *Baseline_2* and *Baseline_3* strategies, as depicted in Fig. 7.

It is evident that with the *FaRe*, the aggregate throughput, regardless of the evaluated *t-Reordering* value, never goes down to zero or reaches abnormal peak values. Indeed, the *FaRe* achieves a stable throughput, which on average is around 25 Mbps for all the cases. Consequently, the application never stops receiving data even though the data splitting via the SN link is temporarily suspended.

It is worth mentioning that the periods with zero throughput and abnormal peak values, i.e., throughput larger than 45 Mbps, appear since the TCP packets received at the transport layer are not in sequence. Hence, they are buffered until TCP recovers the lost packet(s), and they can be delivered in sequence to the upper layers. In our evaluations, the lost TCP packets mainly appear when the PDCP reordering timer expires for a given PDCP PDU. Thus, data with sequence gaps, i.e., missing packets, is received at the transport layer. As shown in Fig. 8, all the Baseline strategies suffer from this problem to a greater or lesser extent, being the *Baseline_1* the strategy in which more spurious PDCP PDU discards occur when the *t-Reordering* expires. Fig. 8 shows the average number of *t-Reordering* timeouts after running the 20 experiments for the *FaRe* and Baseline strategies.

Considering the average aggregate throughput obtained with the evaluated Baseline strategies, i.e., around 24 Mbps, the SN change events have no significant impact on the obtained result regardless of the *t-Reordering* value. However, the instantaneous throughput values significantly vary

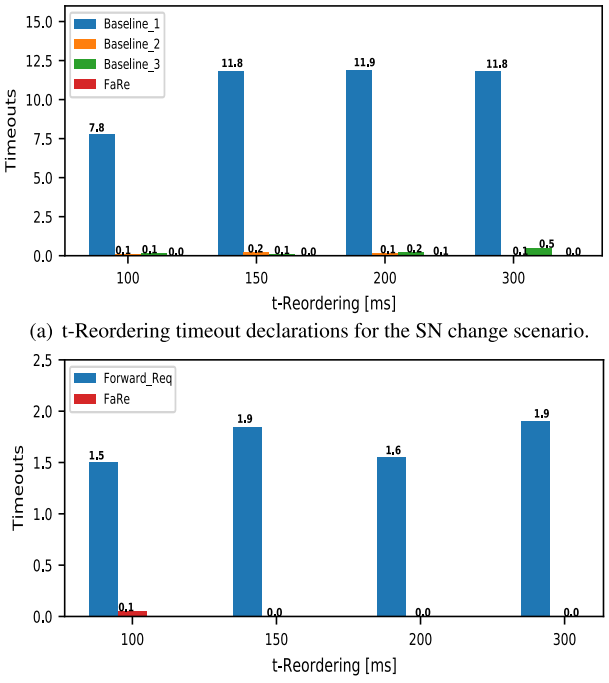

(a) t-Reordering timeout declarations for the SN change scenario.

(b) t-Reordering timeout declarations for the SCG failure scenario.

FIGURE 8. Average number of t-Reordering timeout declarations.

during the events mentioned above. Some applications, such as LL-eMBB or real-time applications, may not tolerate having extended periods of zero throughputs. Hence, it may be more beneficial to have a continuous data flow rather than short periods of zero and very high throughputs. Contrary to the results obtained with the Baseline strategies, the *FaRe* achieves a stable throughput with an average of 25.5 Mbps regardless of the reordering timeout value. As a significant advantage, the FaRe's instantaneous throughputs never reach abnormal zero and very high peak values, as illustrated in Fig. 7.

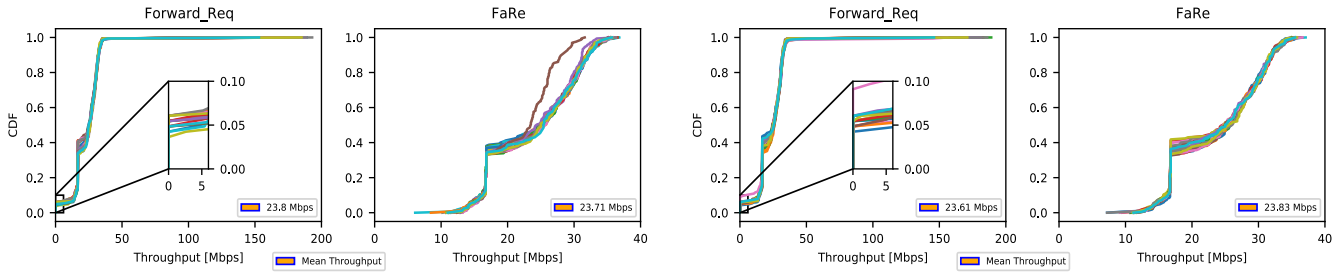

(a) CDF throughput for t-Reordering = 100 ms

(b) CDF throughput for t-Reordering = 300 ms

FIGURE 9. CDF of the obtained aggregate throughput using different reordering timeout values for the SCG failure case.

2) SCG FAILURE SCENARIO

The SCG failure scenario is evaluated in a data session of 20 seconds, in which five SN change events occur before the MN receives the *SCGFailureInformation* message. As detailed in Section V-B2, the SN change events are addressed using the *Baseline_2* strategy. Thus, how the aforementioned strategies manage the SN change and failure events influence the obtained aggregate throughput. Because the UE losses connectivity with the SN after the SCG failure, the UE switches from MR-DC to SC operation, causing the aggregate throughput to drastically drop from theoretical values of 34-35 Mbps to 16-17 Mbps, as observed in Fig. 9.

The aggregate throughput obtained with the *FaRe* is stable and does not have abnormal zero or peak values compared to the *Forward_Req* strategy. The FaRe achieves such stability since it is aware of the last received PDCP PDU and the possible sequence gaps, the information of which is obtained from the PDCP Status Report. Hence, the MN can timely *retransmit* the non-received PDUs and fill the PDCP sequence gaps present at the UE. Since this process is completed within few milliseconds, the performance of the *FaRe* does not depend on the *t-Reordering* value. Similar to the SN change scenario, the average throughput obtained with the *Forward_Req* and the *FaRe* are practically the same. However, some instantaneous throughput values, i.e., the zero and peak values, obtained with the *Forward_Req* strategy can seriously degrade the application's performance.

Additionally, since the UE will no longer receive data via the SN link after the UE declares an *SCG failure*, the probability of having PDCP PDU discards caused by one or more expirations of the *t-Reordering* is high. Because the *Forward_Req* strategy is not efficient in managing such an event, it has a 4-7% probability of obtaining zero throughputs, as depicted in Figs. 9(a) and 9(b) for the cases when the *t-Reordering* is configured with 100 and 300 ms, respectively. As observed in Fig.8(b), with the *Forward_Req* strategy, the average number of timeout declarations and the subsequent effect on the discarded PDCP PDUs and buffering delay, at PDCP and Transport layers, is much larger than for the *FaRe*. The results shown in Figs. 9 and 8(b) reflect the inability of the *Forward_Req* strategy to correctly identify the non-received PDCP PDUs and timely *re-route* them to the UE via the MN's Uu. For this reason, we argue that this method is not suitable to be used with latency- and reliability-constrained applications such as LL-eMBB applications. On the other hand, the *FaRe* has as added value the ability to efficiently serve the aforementioned applications.

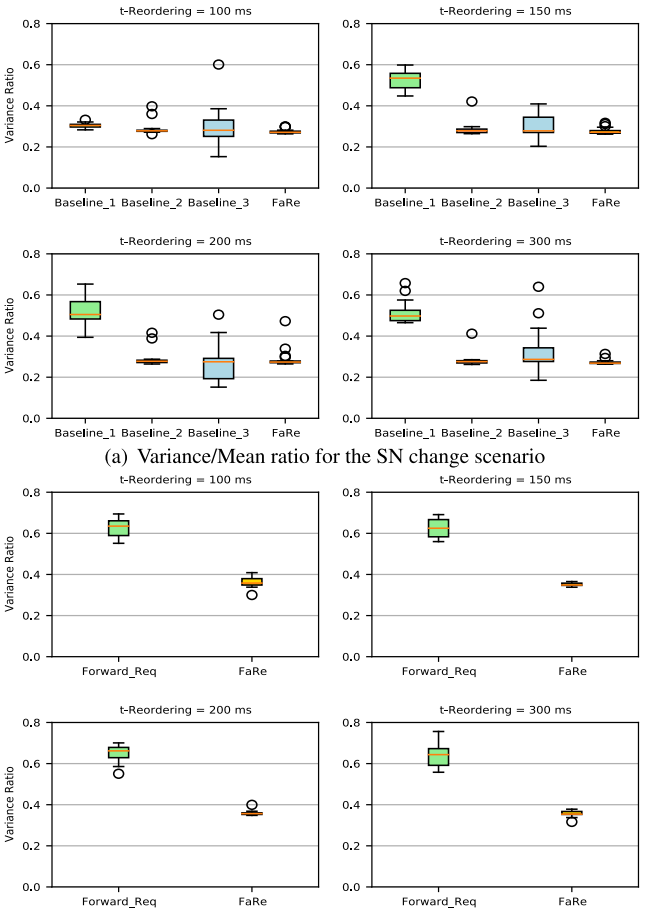

(a) Variance/Mean ratio for the SN change scenario

(b) Variance/Mean ratio for the SCG failure scenario

FIGURE 10. Aggregate throughput variance.

B. THROUGHPUT VARIANCE

In order to study the variability of the obtained aggregate throughput during the entire data session and evince the zero and peak throughput periods, we use the R_{var}. The boxplot graphs presented in Fig. 10 correspond to the results obtained

after running 20 experiments for the analyzed strategies. For this, the boxplot's median represents the average R_{var} obtained for the 20 experiments. Likewise, the first and third quartiles indicate how spread is the R_{var} among the experiments. We observe in Figs. 10(a) and 10(b) that the throughput obtained with the *FaRe* does not present high variability compared to the one obtained with the *Baseline* and *Forward_Req* strategies. The R_{var} for the *FaRe* is on average 0.27 for the SN change scenario and 0.35 for the SCG failure scenario, regardless of the configured *t-Reordering* value. In both cases, the individual R_{var} results are highly concentrated around the average value.

For the SN change scenario, the R_{var} obtained for the *Baseline_1* and *Baseline_3* strategies are widely spread, which means that the throughput is not stable. Moreover, according to the results shown in Fig. 10(a), the *Baseline_2* strategy achieves a stable throughput. Nevertheless, we can see from Fig. 7 that such a strategy still achieves zero and peak throughputs. Furthermore, for the SCG failure scenario, the throughput variance obtained with the *Forward_Req* strategy is on average 85% higher than with the *FaRe*, as shown in Fig. 10(b). This instability is the reason to obtain throughputs that are not concentrated around their average value. Note that the observed throughput instability in the benchmarking strategies is caused by the additional time the received PDUs spend in the PDCP reordering buffer before being delivered to the upper layers. The results presented in Fig. 10 confirm that the aggregate throughput obtained with the *FaRe* is not subject to significant and abnormal variations, which is desirable in real-time applications to satisfy a given Quality of Experience [21].

C. DATA RELIABILITY

During data aggregation, the UE discards all the received PDCP PDUs that do not belong to the same reordering window [9], which typically happens after the *t-Reordering* expires. In this regard, if a PDCP PDU does not arrive at the UE or is discarded, the upper layers will not receive their corresponding data, which may obligate the sender to retransmit the missing data. For this reason, to study the data reliability, we focus on the reliability achieved at the UE's PDCP layer by counting the number of correctly received PDCP PDUs and comparing them with the number of PDUs that left the transmitting PDCP layer. Tables 4 and 5 illustrate the average percentage of PDUs correctly received at the UE and its standard deviation obtained after running 20 experiments with each of the evaluated strategies and reordering timeout values for the *SN change* and *SCG failure* scenarios, respectively.

The results shown in Table 4, for the *SN change* scenario, indicate that the *Baseline_2* and *Baseline_3* strategies cannot achieve reliability of five nines for all the reordering timeout values. At a glance, the obtained reliability seems to be suitable for applications such as LL-eMBB [10]. Nevertheless, such applications require a high and stable data rate [38], which is not possible to achieve with any of the Baseline strategies. On the contrary, the *FaRe* achieves reliability of 100% for all reordering timeout values, except for 300 ms, with the additional benefit of having a stable and smooth data rate, as observed in Figs. 7 and 10(a).

TABLE 4. PDCP reliability for the SN change scenario.

Strategy	100 ms	150 ms	200 ms	300 ms
Baseline_1	99.39029	98.24219	98.20333	98.03971
δ_{B1}	0.02423	0.32628	0.24285	0.10097
Baseline_2	99.99965	99.99957	99.99950	**100**
δ_{B2}	0.00155	0.00187	0.00164	**0**
Baseline_3	99.99943	99.99965	99.99523	99.99831
δ_{B3}	0.00177	0.00154	0.01861	0.00321
FaRe	**100**	**100**	**100**	**99.99981**
δ_{FaRe}	**0**	**0**	**0**	**0.0019**

TABLE 5. PDCP reliability for the SCG failure scenario.

Strategy	100 ms	150 ms	200 ms	300 ms
Forward_Req	99.97903	99.97978	99.97662	99.97464
δ_{Fr}	0.02771	0.02355	0.02370	0.03080
FaRe	**99.99973**	**100**	**100**	**100**
δ_{FaRe}	**0.00154**	**0**	**0**	**0**

Furthermore, for the *SCG failure* scenario, the results shown in Table 5 demonstrate the impossibility of the *Forward_Req* strategy to offer the reliability of more than three nines, which is not the case for the *FaRe*. Indeed, with the *FaRe*, more than the 99.999% of the PDUs that left the MN's PDCP layer are correctly received at the UE, even though the SN link suddenly fails.

D. DATA INTERRUPTION TIME

To study the data interruption time, we rely on the periodic throughput reports delivered by the *iperf3* tool. In this regard, we measure the periods of zero throughputs during the entire data session for each experiment. For this, the throughput reports are periodically collected in our experiments every 100 ms for both scenarios. The results shown in Fig. 11 represent the average data interruption time experienced by the transport layer after running 20 experiments for the *FaRe* and every benchmarking strategy.

The results depicted in Fig. 11(a), for the *SN change* scenario, show the effectiveness of the *FaRe* to avoid suffering from data interruption periods. On the other hand, it can be visualized that *Baseline_2* creates the lowest average interruption time among the Baseline strategies. At a glance, the obtained interruption time may not represent a significant problem in scenarios where the throughput stability and the data reliability are not the primary concern. However, for latency- and reliability-constrained applications, the *Baseline_2* is not an option to consider. Even though the Baseline strategies and the *FaRe* achieve on average a similar aggregate throughput, as shown in Fig. 7, the *SN change*

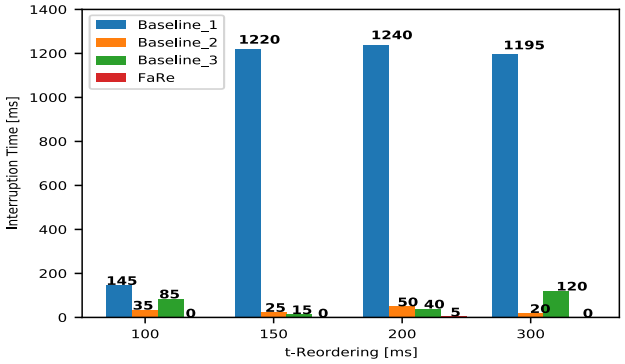

(a) Data interruption time for the SN change scenario

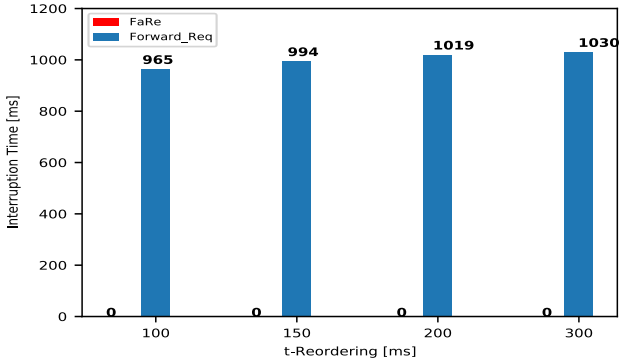

(b) Data interruption time for the SCG failure scenario

FIGURE 11. Data interruption time at the transport layer level.

events are not efficiently managed by the Baseline strategies. Therefore, it will be challenging for applications to meet their required KPIs.

For the *SCG failure* scenario, the results obtained with the *FaRe* and illustrated in Fig. 11(b) demonstrate that it is possible to maintain a continuous data flow to the UE, even though the SN link fails. The novel capability of the *FaRe* to retransmit the non-received PDUs makes the application never stop receiving data, which on the contrary, significantly differs from the results obtained with the *Forward_Req* strategy. It is important to remark that the UE switches to SC operation after the *SCG failure* event, so the UE keeps receiving data via the MN link. Consequently, if the missing PDCP PDU(s) do not arrive at the UE in time to fill the PDCP sequence gap, the transport layer will not receive the expected data. For reliability-oriented protocols such as TCP, the transport layer will retransmit the missing data, significantly reducing the throughput and increasing the data interruption periods.

VII. THE *FaRe* IMPLEMENTATION IMPACT

In Section VI, we have demonstrated the effectiveness of the *FaRe* to reduce the application's data interruption time during *SN change* and *SCG failure* events. The results show that it is possible to get a stable aggregate throughput with high data reliability thanks to the *FaRe*. These benefits are obtained since the MN temporarily stores a copy of the PDCP PDUs split via the SN and timely retransmits them when one of the events described above occurs. Since, in our experiments, the

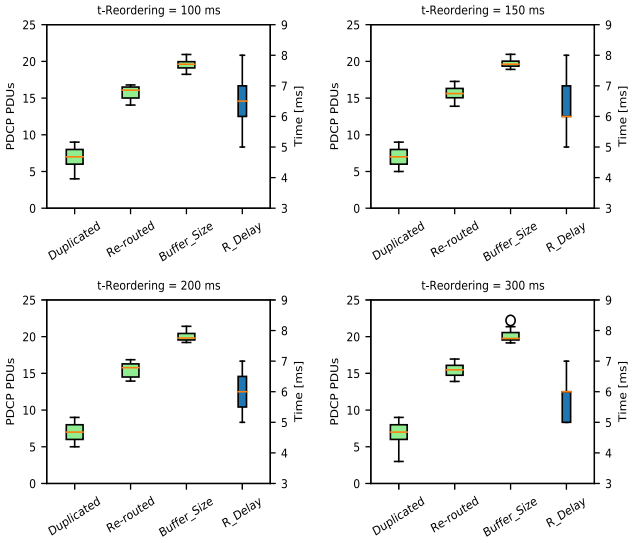

FIGURE 12. Different statistics for the implementation impact of the *FaRe*.

FaRe-PDUs are acknowledged by the MN every 5 ms, the memory allocation requirements at the MN are negligible. In fact, the *FaRe-Buffer* usage results, depicted as *Buffer_Size* in Fig. 12, show that during a data session of 20 seconds, on average, 20 PDUs are present in the *FaRe-Buffer*. In our LTE/NR testbed, the *iperf3* tool generates packets of fixed size, creating PDCP PDUs of 1466 bytes. Hence, the average *FaRe-Buffer* size corresponds to 29.3 KBytes, which has a negligible impact on the overall MN's performance. It is worth mentioning that the buffer demand may slightly increase in case of higher throughput demands and larger bandwidth sizes.

Moreover, during an *SCG failure* event, the *FaRe retransmits* via the MN's Uu, on average, 15 *FaRe-PDUs*, as can be visualized with the variable *Re-routed* in Fig. 12. The first of these PDUs arrive at the UE's PDCP in approximately 5-8 ms. This delay, shown as *R_Delay* in Fig. 12, matches with the theoretical delay, i.e., Δ^*_{FaRe}, computed with (2). Note that the measurement of this delay starts once the MN receives the *SCGFailureInformation* message and ends when the UE's PDCP layer receives the first *FaRe-PDU*.

Additionally, we noted in our experiments that in some *SN change* events, some PDUs, that were initially transmitted via the SN, arrive at the UE's PDCP after the *FaRe-PDUs*. This event happens when several HARQ retransmissions were required in the UE-SN path to decode the transport block correctly. During the evaluated period, the UE received, on average, 6 duplicated PDUs, as depicted with the variable name *Duplicated* in Fig. 12. This random event has a negligible impact on the performance of the *FaRe* since the throughput, reliability, and data interruption are not affected.

VIII. CONCLUSION

A fast data recovery mechanism that minimizes the data interruption time experienced by the application in MR-DC scenarios with mobility is presented in this article. The

proposed mechanism, named FaRe, intelligently identifies and timely retransmits, via the MN, the PDCP PDUs that were not received at the UE or transmitted by the SN during the aforesaid events. To accomplish this, the MN stores at its PDCP layer a copy of every PDCP PDU split via the SN until an acknowledgment is received. Hence, when an SN change or SCG failure occurs, the non-delivered or non-transmitted PDCP PDUs can timely be transferred to the UE via the MN's Uu, thus, avoiding time-consuming higher layer retransmissions or slow data forwarding procedures.

The experimental evaluations conducted on an LTE/NR testbed, build using the OAI software, revealed that when the UE aggregates data, the SN change or SCG failure events affect the throughput stability and latency requirements of TCP-based applications. The results obtained in the evaluated setup showed that the proposed *FaRe* mechanism effectively reduces the harmful effects of such events for all the PDCP reordering timeout assessed values. Indeed, the *FaRe* results in zero data interruption periods for all the assessed scenarios except when the *t-Reordering* is 200 ms in the SN change scenario. The efficient approach used by the *FaRe* allows the UE to achieve a stable and high aggregate throughput with data reliability of more than 99.999% without upper layer retransmissions. This is a desirable feature to support emerging 5G applications such as low latency eMBB.

On the contrary, with the alternative strategies used to address these events, i.e., the *Baseline* and *Forward_Req* strategies, the application stops receiving data for periods ranging from dozens to hundreds of milliseconds. These long data interruption periods make the PDCP reordering mechanism discard several PDCP PDUs, diminishing the data reliability to three nines in the *SCG failure* scenario and even to two nines in the *SN change* scenario, with the *Baseline_1* strategy. For this reason, it is challenging to meet the reliability and latency targets defined for some 5G use cases.

Finally, we have demonstrated that the novel inclusion of the PDCP Status Report within the *SCGFailureInformation* message helps the proposed *FaRe* mechanism to effectively reduce the application's data interruption time caused by failures in the SN. Since, in a typical MR-DC deployment, the UE may experience HOs and failures at the MN and SN, we plan to extend the functionality of the *FaRe* to manage the Inter-Master Node HOs and MCG failures as well. Thus, assuring seamless network connectivity via both BSs. In addition, we believe that the *FaRe* is a reliable solution that can be used to avoid missing data during SCG activation/deactivation procedures. Yet, it is still necessary to study solutions that allow such integration.

REFERENCES

[1] *Universal Mobile Telecommunications System (UMTS); LTE; 5G; NR; Multi-Connectivity; Overall Description; Stage-2*, document TS 37.340, V17.0.0, 3GPP, 2022.
[2] C. Pupiales, D. Laselva, Q. De Coninck, A. Jain, and I. Demirkol, "Multi-connectivity in mobile networks: Challenges and benefits," *IEEE Commun. Mag.*, vol. 59, no. 11, pp. 116–122, Nov. 2021.
[3] C. Pupiales, D. Laselva, and I. Demirkol, "Capacity and congestion aware flow control mechanism for efficient traffic aggregation in multi-radio dual connectivity," *IEEE Access*, vol. 9, pp. 114929–114944, 2021.
[4] C. Rosa, K. Pedersen, H. Wang, P. H. Michaelsen, S. Barbera, E. Malkaméki, T. Henttonen, and B. Sébire, "Dual connectivity for LTE small cell evolution: Functionality and performance aspects," *IEEE Commun. Mag.*, vol. 54, no. 6, pp. 137–143, Jun. 2016.
[5] O. N. C. Yilmaz, O. Teyeb, and A. Orsino, "Overview of LTE-NR dual connectivity," *IEEE Commun. Mag.*, vol. 57, no. 6, pp. 138–144, Dec. 2019.
[6] R. Dilli, "Analysis of 5G wireless systems in FR1 and FR2 frequency bands," in *Proc. 2nd Int. Conf. Innov. Mech. Ind. Appl. (ICIMIA)*, Mar. 2020, pp. 767–772.
[7] *ETRI*, document R2-1903906, 3GPP, Xi'an, China, 2019.
[8] Ericsson. *Tackling Fast Recovery from Radio Link Failure—Ericsson*. Accessed: Mar. 25, 2022. [Online]. Available: https://www.ericsson.com/en/blog/2020/9/fast-recovery-from-radio-link-failure
[9] *5G; NR; Packet Data Convergence Protocol (PDCP) Specification*, document TS 38.323, V17.0.0, 3GPP, 2022, pp. 0–27.
[10] *MediaTek*, document RWS-210094, MediaTek, Hsinchu, Taiwan, 2021.
[11] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi, "Will TCP Work in mmWave 5G cellular networks?" *IEEE Commun. Mag.*, vol. 57, no. 1, pp. 65–71, Jan. 2019.
[12] *5G; NR; Radio Resource Control (RRC); Protocol Specification*, document TS 38.331, V17.0.0, 3GPP, 2022.
[13] *Ericsson*, document R2-170271, Ericsson, Spokane, DC, USA, 2017.
[14] B. Jin, S. Kim, D. Yun, H. Lee, W. Kim, and Y. Yi, "Aggregating LTE and Wi-Fi: Toward intra-cell fairness and high TCP performance," *IEEE Trans. Wireless Commun.*, vol. 16, no. 10, pp. 6295–6308, Oct. 2017.
[15] K. Nguyen, M. G. Kibria, J. Hui, K. Ishizu, and F. Kojima, "Minimum latency and optimal traffic partition in 5G small cell networks," in *Proc. IEEE Veh. Technol. Conf.*, Jun. 2018, pp. 1–5.
[16] D. Lopez-Perez, D. Laselva, E. Wallmeier, P. Purovesi, P. Lunden, E. Virtej, P. Lechowicz, E. Malkamaki, and M. Ding, "Long term evolution-wireless local area network aggregation flow control," *IEEE Access*, vol. 4, pp. 9860–9869, 2016.
[17] H. Wang, C. Rosa, and K. I. Pedersen, "Dual connectivity for LTE-advanced heterogeneous networks," *Wireless Netw.*, vol. 22, no. 4, pp. 1315–1328, 2016.
[18] M.-S. Pan, T. M. Lin, C. Y. Chiu, and C. Y. Wang, "Downlink traffic scheduling for LTE—A small cell networks with dual connectivity enhancement," *IEEE Commun. Lett.*, vol. 20, no. 4, pp. 796–799, Apr. 2016.
[19] R. P. Antonioli, E. B. Rodrigues, D. A. Sousa, I. M. Guerreiro, C. F. M. E. Silva, and F. R. P. Cavalcanti, "Adaptive bearer split control for 5G multi-RAT scenarios with dual connectivity," *Comput. Netw.*, vol. 161, pp. 183–196, Oct. 2019.
[20] R. P. Antonioli, I. M. Guerreiro, D. A. Sousa, E. B. Rodrigues, C. F. M. E. Silva, T. F. Maciel, and F. R. P. Cavalcanti, "User-assisted bearer split control for dual connectivity in multi-RAT 5G networks," *Wireless Netw.*, vol. 26, no. 5, pp. 3675–3685, Jul. 2020.
[21] M. Polese, M. Giordani, M. Mezzavilla, S. Rangan, and M. Zorzi, "Improved handover through dual connectivity in 5G mmWave mobile networks," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 9, pp. 2069–2084, Sep. 2017.
[22] P.-J. Hsieh, W.-S. Lin, K.-H. Lin, and H.-Y. Wei, "Dual-connectivity prevenient handover scheme in control/user-plane split networks," *IEEE Trans. Veh. Technol.*, vol. 67, no. 4, pp. 3545–3560, Apr. 2018.
[23] T. Mumtaz, S. Muhammad, M. I. Aslam, and N. Mohammad, "Dual connectivity-based mobility management and data split mechanism in 4G/5G cellular networks," *IEEE Access*, vol. 8, pp. 86495–86509, 2020.
[24] R. El Banna, H. M. ELAttar, and M. Aboul-Dahab, "Handover scheme for 5G communications on high speed trains," in *Proc. 5th Int. Conf. Fog Mobile Edge Comput. (FMEC)*, Apr. 2020, pp. 143–149.
[25] K. Qi, T. Liu, C. Yang, S. Suo, and Y. Huang, "Dual connectivity-aided proactive handover and resource reservation for mobile users," *IEEE Access*, vol. 9, pp. 36100–36113, 2021.
[26] L. C. Gimenez, P. H. Michaelsen, K. I. Pedersen, T. E. Kolding, and H. C. Nguyen, "Towards zero data interruption time with enhanced synchronous handover," in *Proc. IEEE 85th Veh. Technol. Conf. (VTC Spring)*, Jun. 2017, pp. 1–6.
[27] *LTE; Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall Description; Stage 2*, document TS 36.300, V17.0.0, 3GPP, 2022.
[28] I. Viering, H. Martikainen, A. Lobinger, and B. Wegmann, "Zero-zero mobility: Intra-frequency handovers with zero interruption and zero failures," *IEEE Netw.*, vol. 32, no. 2, pp. 48–54, Mar./Apr. 2018.

[29] H. Martikainen, I. Viering, A. Lobinger, and T. Jokela, "On the basics of conditional handover for 5G mobility," in *Proc. IEEE 29th Annu. Int. Symp. Pers., Indoor Mobile Radio Commun. (PIMRC)*, Sep. 2018, pp. 1–7.

[30] *5G; NR; NR and NG-RAN Overall Description*, document TS 38.300, V17.0.0, 3GPP, 2022.

[31] *5G; Study on Scenarios and Requirements for Next Generation Access Technologies*, document TR 38.913, V17.0.0, 3GPP, 2022, pp. 1–42.

[32] *Study on Latency Reduction Techniques for LTE*, document TR 36.881, V14.0.0, 3GPP, 2016.

[33] *Universal Mobile Telecommunications System (UMTS); LTE; Virtual Reality (VR) Media Services Over 3GPP*, document TR 26.918, V16.0.0, 3GPP, 2020.

[34] M. Gerasimenko, D. Moltchanov, M. Gapeyenko, S. Andreev, and Y. Koucheryavy, "Capacity of multiconnectivity mmWave systems with dynamic blockage and directional antennas," *IEEE Trans. Veh. Technol.*, vol. 68, no. 4, pp. 3534–3549, Apr. 2019.

[35] V. Petrov, D. Solomitckii, A. Samuylov, M. A. Lema, M. Gapeyenko, D. Moltchanov, S. Andreev, V. Naumov, K. Samouylov, M. Dohler, and Y. Koucheryavy, "Dynamic multi-connectivity performance in ultra-dense urban mmWave deployments," *IEEE J. Sel. Areas Commun.*, vol. 35, no. 9, pp. 2038–2055, Sep. 2017.

[36] M. F. Ozkoc, A. Koutsaftis, R. Kumar, P. Liu, and S. S. Panwar, "The impact of multi-connectivity and handover constraints on millimeter wave and terahertz cellular networks," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 6, pp. 1833–1853, Jun. 2021.

[37] S. Choi, J.-G. Choi, and S. Bahk, "Mobility-aware analysis of millimeter wave communication systems with blockages," *IEEE Trans. Veh. Technol.*, vol. 69, no. 6, pp. 5901–5912, Jun. 2020.

[38] S. Jun, Y.-S. Choi, and H. Chung, "Considerations on ultra broadband, high reliable and low latency services in 6G system," in *Proc. Int. Conf. Inf. Commun. Technol. Converg. (ICTC)*, Oct. 2021, pp. 1552–1554.

[39] *ETRI*, document R2-1903906, MIT in eMBB Handover, 3GPP, Xi'an, China, 2019.

[40] *Docomo*, document RWS-210234, Docomo, Tokyo, Japan, 2021.

[41] *Samsung*, document RWS-210183, Samsung, Suwon-si, (South) Korea, 2021.

[42] D. Kaspar, "Multipath aggregation of heterogeneous access networks," Ph.D. dissertation, Dept. Math. Natural Sci., Univ. Oslo, Oslo, Norway, 2012.

[43] *5G; NG-RAN; NR User Plane Protocol*, document TS 38.425, V17.0.0, 3GPP, 2022.

[44] N. Nikaein, M. K. Marina, S. Manickam, A. Dawson, R. Knopp, and C. Bonnet, "OpenAirInterface: A flexible platform for 5G research," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 5, pp. 33–38, Oct. 2014.

[45] C. Pupiales, W. Nitzold, C. Felber, and I. Demirkol, "Software-based implementation of dual connectivity for LTE," in *Proc. IEEE 16th Int. Conf. Mobile Ad Hoc Sensor Syst. Workshops (MASSW)*, Nov. 2019, pp. 178–179.

[46] *Open-Access Research Testbed for Next-Generation Wireless Networks (ORBIT)*. Accessed: May 3, 2021. [Online]. Available: https://www.orbit-lab.org/

[47] *Study on New Radio Access Technology Physical Layer Aspects*, document TR 38.802, V14.2.0, 3GPP, 2017.

[48] F. Kaltenberger, "OpenAirInterface 5G overview, installation, usage," in *OpenAirInterface Workshop*. Beijing, China: OpenAirInterface Alliance, 2019.

[49] *LTE; Scenarios and Requirements for Small Cell Enhancements for E-UTRA and E-UTRAN*, document TR 36.932, V12.1.0, 3GPP, 2014, pp. 0–15.

CARLOS PUPIALES (Member, IEEE) received the B.Sc. degree in telecommunications engineering from the Escuela Politécnica Nacional, Ecuador, and the M.Sc. degree in telecommunications engineering from the University of Melbourne, Australia. He is currently pursuing the Ph.D. degree with the Department of Network Engineering, Universitat Politècnica de Catalunya, Spain. His research interest includes the area of network protocols in mobile networks. He was a recipient of the Best Demo Award in IEEE MASS 2019.

DANIELA LASELVA received the M.Sc. degree in electrical engineering from the Polytechnic of Bari, Italy, in 2002. From 2002 to 2006, she was active in COST273 and EU FP6 Project WINNER on MIMO radio channel modeling with Elektrobit, Finland, and worked as a Senior Design Engineer with Nokia, Finland. Since 2006, she has been with Nokia, Aalborg, Denmark. She is currently with the Nokia Standardization and Research Team as a Senior Research Specialist. She is currently engaged with 5G new radio design and standardization, including solutions for efficient transmission of small data, UE power saving, multi-connectivity, and support of mission-critical applications. She has coauthored more than 30 peer-reviewed publications and four book chapters, and she is the inventor of numerous patents on a wide range of topics.

ILKER DEMIRKOL (Senior Member, IEEE) received the B.Sc., M.Sc., and Ph.D. degrees in computer engineering from Boğaziçi University, İstanbul, Turkey. He is currently an Associate Professor with the Department of Mining, Industrial and ICT Engineering, Universitat Politècnica de Catalunya, where he works on wireless networks and wake-up radio systems. His research interests include communication protocol development for the aforementioned networks, along with performance evaluation and optimization of such systems. He was a recipient of the 2010 Best Mentor Award from the Electrical and Computer Engineering Department, University of Rochester, Rochester, NY, USA.