Evaluating Handover Performance for End-to-End LTE Networks with OpenAirInterface

Rodolphe Bertolini and Mickael Maman
CEA-Leti, Université Grenoble Alpes, F-38000 Grenoble, France
{rodolphe.bertolini, mickael.maman}@cea.fr

***Abstract*—This paper presents an experimental testbed of the handover procedure using an accessible and reconfigurable software defined radio environment with an end-to-end architecture (i.e. including both the radio access network and the core network). First, we provide a comprehensive overview of the X2 handover procedure, in an end-to-end cellular network architecture and detail the handover condition, message flow and latency decomposition. We then describe our OpenAirInterface based implementation and end-to-end experimentation of LTE X2 handover in a full-SDR environment. Finally, we analyze the performance in terms of end-to-end throughput and of latency for each step of the handover procedure and compare it to the state of the art of X2 handover experimentation.**

I. Introduction

Handover (HO) in cellular network is a mechanism that provides continuity of service to a terminal with varying radio or traffic quality by switching its connection from the serving cell to a selected neighboring cell. In regular 4G Long-Term Evolution (LTE) / 5G New Radio (NR) networks, the HO decision is made by the serving eNB, with knowledge of the different signal strengths that the connected terminal detects. The HO mechanism is inherent in contexts such as a vehicle-to-everything environment due to high the mobility of vehicles or in the factory of the future due to the blockage of moving objects, the mobility of a terminal and the prioritization of one link over another (e.g., if a specific terminal has a specific quality of service to achieve). In [1], M. Tayyab et al. did a survey on innovative LTE and 5G HO techniques but this comprehensive knowledge base does not provide any experimentation reference. In general, HO procedures clearly lack the means to experiment with them. In [2], Han et al. evaluate the performance of X2 HO (as explain in section II.A) in real world environment. They use a Commercial Off The Shelf (COTS) User Equipment (UE) such as smartphone, connected to wireshark to dissect the control messages and can track the message flow. Evaluating performance in a real world environment provides insight into where the latency bottleneck is. However, since they do not have access to the core network, most of the data they bring in on delays in the core network is inferred from procedures that are not related to the HO.

Software Defined Radio (SDR) is a paradigm that enables flexible radio systems. Indeed, in SDR systems, most of the processing is delegated to software computation (e.g., turbo code encoder/decoder [3]). In the hardware remain the procedures related to the radio (e.g., transmission, reception and ADC/DAC). OpenAirInterface [4] (OAI) is an open-source framework that aims to provide a pluggable cellular network solution using SDR boards. It implements Radio Area Network (RAN) elements, consisting of the UE and the eNodeB (eNB). On the Core Network (CN) side, OAI implements the Mobility Management Entity (MME), the Home Subscriber Server (HSS), and the Serving Gateway and Packet Data Network Gateway that are combined into a single entity (SPGW). However the Policy and Charging Rules Function (PCRF) is not implemented. OAI currently supports 4G LTE, and non standalone 5G NR limited to the use of COTS UE. In [5], Alexandris et al. present the HO in OAI, but they do not specify which layers are enabled or not. Similarly, they do not mention the presence of a CN in their experimental setup. In [6], Manco et al. are experimenting LTE V2X with OAI, using sidelink capabilities. However, there is no mention of mobility which is a major point in V2X as vehicles are more likely to travel distances not covered by a single cell.

To our knowledge, there is no experiment with the HO procedure using an accessible and reconfigurable environment with an end-to-end architecture (i.e. including both RAN and CN). An implementation of HO in OAI exists but only with COTS UE. The use of a COTS UE reduces the degree of freedom of research experimentation. Indeed, we may be limited by the implementation of standards and constructors. Having a full SDR experimentation setup allows us to bypass these limitations. For instance, if we want to evaluate a scenario in which the eNB or another entity has full control over the HO trigger and decision, we can disable the UE measurement reports to avoid unnecessary control traffic, which is not possible with a COTS UE. In this paper, we address an end-to-end, configurable experimental testbed for HO procedure. The testbed is end-to-end as it involves components from a UE to the SPGW CN and is configurable as we have access to the code of any RAN and CN component and can customize it. Our contribution is threefold: first, we implemented LTE X2 HO in the OAI RAN. This includes (but is not limited to) the implementation of multi-eNB management, such as eNB synchronization eNB scanning and selection by UE RAN layers, Reference Signal Received Power (RSRP) measurement of neighboring cells, Contention-Free Random Access (CFRA) versus Contention-Based Random Access (CBRA) procedure. Next, we experimented with X2 HO for realistic mobility scenarios in an end-to-end, full-SDR environment with an accurate channel emulator. Finally, we analyzed the performance in terms of end-to-end throughput and latency for each step of the procedure.

The paper is organized as follows. Section II gives an overview of LTE architecture and HO procedure. Section III presents our experimental testbed and shows the performance of handover procedure. Section IV concludes the paper.

II. LTE architecture and HO flow

In this section, we recall the LTE architecture and the Random Access (RA) procedure. We also provide an overview of the HO procedure we implement, including the HO message flow, HO conditions and HO latency decomposition.

A. LTE Architecture

LTE is split into two entities as illustrated in Figure 1: the evolved universal terrestrial radio access network which we refer to as RAN in this paper, and the Evolved Packet Core (EPC), also referred to as CN. RAN is composed of UEs and eNBs. The link between the eNBs is called X2.

When the UE is switched on, it searches for surrounding eNBs and selects one. It establishes the connection with the eNB by performing a RA procedure via the Physical Random Access Channel (PRACH). There are two RA procedures: Figures 2 and 3 describe the CBRA and the CFRA procedures, respectively. In CBRA, the UE chooses a random number that is used by the eNB to identify the new UE. It involves additional control, as several UEs may choose the same RA

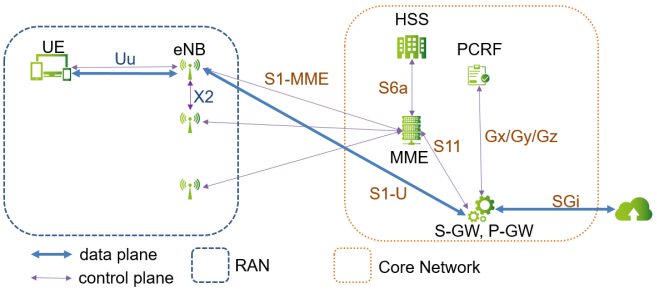

Fig. 1. LTE end to end architecture

preamble. In CFRA, the preamble is given by the eNB through the previous signaling. The RA in the HO can be either one or the other. In our testbed, we choose the CBRA procedure for the initial connection and the CFRA procedure for the HO because it involves less control traffic, thus reducing the HO time. The connection procedure ends with a *Radio Resource Control (RRC) connection reconfiguration* message sent by the eNB to the UE. Amongst many pieces of information, it contains neighboring cell identifiers used by the UE to measure the Cell-specific References Signals (CRS).

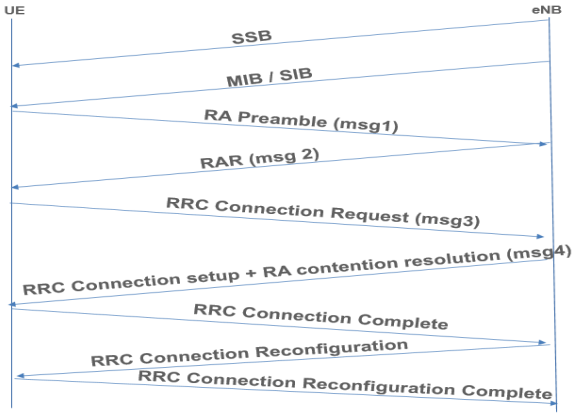

Fig. 2. Message flow of contention-based RA procedure

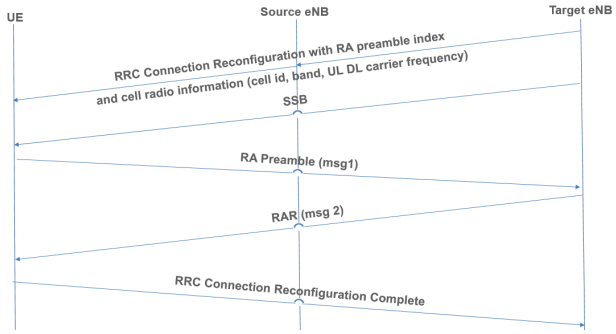

Fig. 3. Message flow of contention-free RA procedure

Once this procedure is complete, the eNB asks the SPGW to establish a data bearer, then the UE is connected to the CN, and the eNB regularly requests information with *Downlink Control Information* (DCI) messages. For example, it asks the UE for the status of its uplink buffer, and the UE replies by sending a *Buffer Status Report* (BSR) message. This exchange is referred to as "DCI to ULSCH traffic" in Figure 5. The downlink (DL) is transferred from the SPGW to the eNB, which sends it to the UE ("1" and "2" in Figure 5).

The eNBs are connected to the CN through the S1 interface. The CN is composed of MME and HSS that manage mobility and user connections, SPGW that routes user traffic and user control, and PCRF that manages the plan and billing.

B. Handover

As stated in the introduction, HO is a mechanism that provides service continuity for a UE whose connection needs to be switched from the serving to a chosen neighboring cell.

1) Handover Condition: Once per subframe (1ms), the UE analyzes the CRS of the serving cell and neighboring cells to estimate their Reference Signal Received Power (RSRP). If the RSRP of the serving cell and neighboring cells verify the input condition (Equation 1) during the Time To Trigger (TTT) without verifying the output condition (Equation 2), the UE sends a measurement report to the serving eNB. This report includes the power values and the ID of the reported cell.

$$Mn + Ofn + Ocn - Hys > Ms + Ofs + Ocs + Off \quad (1)$$

$$Mn + Ofn + Ocn + Hys < Ms + Ofs + Ocs + Off \quad (2)$$

where Mn, Ms are measurement (i.e. RSRP in dBm) of the neighboring or the serving cell, respectively, Ofn, Ofs are frequency-specific offset, depending on the frequency of the neighboring cell or the serving cell, respectively (in dB), Ocn, Ocs are cell-specific offset of the neighboring cell or the serving cell, respectively (in dB), Hys is hysteresis parameter (in dB) and Off is an event offset (in dB).

2) X2 / S1 Handover: LTE defines two types of HO: X2 HO and S1 HO, which are shown in Figure 4.

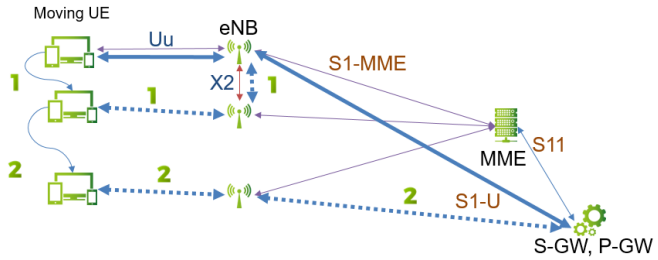

Fig. 4. Data plane path for LTE end to end architecture during X2 HO (1) and intra-MME S1 HO (2)

In X2 HO, DL data from the CN to the UE continues to be sent to the source eNB, which forwards it to target eNB via X2 (Data Plane 1 in Figure 4). Once the UE is connected to the target eNB, a path change is requested by the target eNB to the MME, so the DL traffic is directly forwarded to the target eNB. S1 HO occurs when the MME serving the source eNB and the target eNB are different, or when the X2 link is unavailable. In such case, there is no path switch. The CN creates a new bearer, and the DL data is buffered to the target eNB via S1 link through the newly created bearer. In both cases, after the completion of HO, the data that is buffered at the target eNB during HO is sent to the UE before any other data. In this paper, we consider X2 HO.

3) Handover Message flow: Figure 5 details the X2 HO message flow. Initially, the X2 link is established between the eNBs and the UE connects to the source eNB.Then, the UE moves away from the source eNB to the target eNB. It meets the HO conditions and sends a measurement report to the serving eNB. The source eNB requests, via the X2 interface, the target eNB to perform the HO with a *HO request* message. If the target eNB has enough resources (radio, computing, scheduling), it responds with an *HO request acknowledgement* message. It also sends a *RRC connection reconfiguration* message, which the source eNB forwards to the UE. This RRC message includes a random access preamble and a mobility control information section. With the presence of this section, the UE knows that it must perform an HO. It

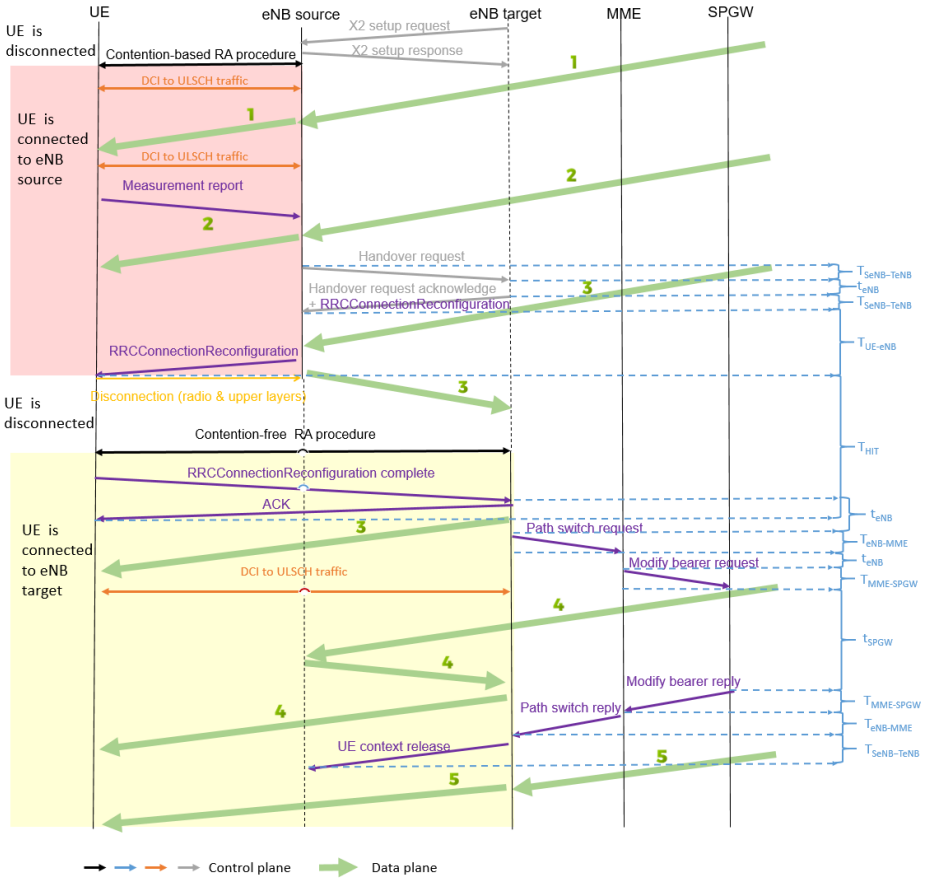

Fig. 5. Control and data message flow before, during and after an X2 handover

reconfigures its layers according to the previous RRC message and disconnects from the source eNB. From this point on, the UE cannot receive or send any information. Since this is an X2 HO, the DL data destined for the UE is still sent by the SPGW to the source eNB, which redirects it to the target eNB ("3" and "4" in Figure 5) until bearer path change is completed. The UE looks for the synchronisation signal from the target eNB, and then performs a CFRA procedure. As soon as it receives its identifier from the target eNB, it replies with a *RRC reconfiguration complete* message and the target eNB requests to the SPGW to change the bearer path. When the change is made, the source eNB releases its UE context and no longer receives DL data addressed to the UE: The CN sends DL data directly to the target eNB ("5" in Figure 5).

4) Handover Latency Decomposition: Han et al. [2] propose Equation 3 which expresses the total HO time T_{HO}. This equation is composed of terms according to the three HO phases : preparation, execution and completion.

$$T_{HO} = T_{HOPrep} + T_{HOExe} + T_{HOComp} \quad (3)$$

The HO preparation time T_{HOPrep} is from the time the source eNB receives measurement the report to the time the UE receives the *RRC connection reconfiguration* message.

$$T_{HOPrep} = 2T_{SeNB-TeNB} + t_{eNB} \quad (4)$$

Where $T_{SeNB-TeNB}$ is the latency experienced through the X2 link and t_{eNB} is the processing time at the eNB.

The HO execution time T_{HOExe} is between the end of the preparation phase and the moment when the UE receives the *RRC connection reconfiguration complete acknowledgement*.

$$T_{HOExe} = T_{HIT} + T_{UE-eNB} \quad (5)$$

Where T_{HIT} is the time between receiving the *RRC connection reconfiguration* message and receiving the ACK of *RRC connection reconfiguration complete* (UE receives ACK from eNB) and T_{UE-eNB} is the time taken to transmit *RC connection reconfiguration* (eNB to UE).

The HO completion time T_{HOComp} is from the time the target eNB receives RRC connection reconfiguration complete message until the source eNB releases the UE context.

$$\begin{aligned} T_{HOComp} =& 2T_{eNB-MME} + 2T_{MME-PGW} + T_{IP-CAN} \\ &+ T_{SeNB-TeNB} + t_{SPGW} + t_{MME} + t_{eNB} \end{aligned} \quad (6)$$

Where $T_{eNB-MME}$ is the latency experienced by the S1-MME link, $T_{MME-PGW}$ is the latency experienced by the S11 link, T_{IP-CAN} is the time needed to change the bearer at the CN side (SPGW and PCRF), $T_{SeNB-TeNB}$ is the latency experienced by the X2 link, t_{SPGW} is the processing time at the SPGW, t_{MME} is the processing time at the MME and t_{eNB} is the processing time at the eNB.

In our testbed, OAI does not include a PCRF. Thus there is no T_{IP-CAN} processing. Moreover, the target eNB sends the RRC connection reconfiguration complete ACK in parallel with the path change request. This parallelization of procedures simplifies Equation 3 as follows:

$$\begin{aligned} T_{HO} =& 3T_{SeNB-TeNB} + 2T_{MME-SPGW} + 2T_{eNB-MME} \\ &+ T_{HIT} + t_{MME} + t_{SPGW} + 2t_{eNB} \end{aligned} \quad (7)$$

III. Experimentation

A. Experimental Testbed

In this section, we present our experimental setup as shown in Figure 6. Our testbed consists of 3 USRPs connected to 3

high-end computers. The target eNB is composed of a USRP x310 in scenario 1, a USRP b210 in scenario 2, connected to a laptop (green rectangles in Figure 6) while the source eNB is a USRP b210 connected to another laptop (red rectangles). The UE consists of a USRP b210 connected to a desktop machine (blue rectangles). This computer also hosts a virtual machine that runs the OAI EPC (pink rectangle). This configuration can represent a MEC-enabled network architecture, or a private network in a factory where the CN is close to RAN.

Since 2.68GHz belongs to the licensed band 7 in Europe, we use SMA cables instead of antennas to interconnect USRPs. As recommended by Ettus for loopback configurations, in scenario 1 the target eNB and the UE have a static 30dB attenuation at their respective Tx. Table I summarizes the LTE network parameters.

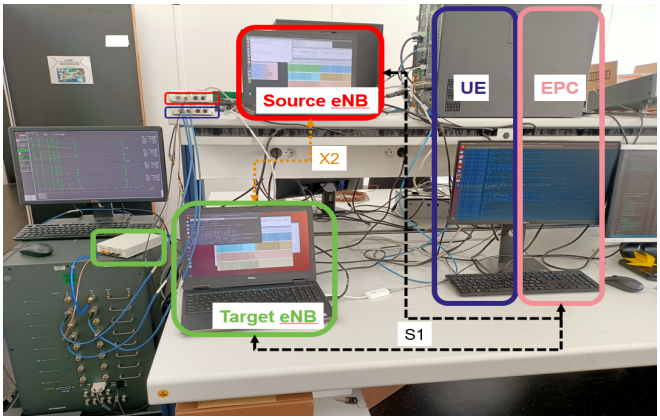

Fig. 6. OAI-based experimental setup for scenario 2

TABLE I
EXPERIMENTATION PARAMETERS

LTE Parameters	**Value**
FDD/TDD	FDD
Downlink Central Frequency	2.68 GHz
Bandwidth	5MHz
Source eNB output power at CF	-60 dBm
Attenuator at Source eNB Tx *	Variable (0-122 dB)
Target eNB output power at CF	-65 dBm
Attenuator at Target eNB Tx*	40 dB
PHICH	1/6
PRACH Configuration Index	0
Max RACH TX	10
RACH Power Ramping Step	4
Handover Parameters	**Value**
Hys	2 dBm
Off = Ofn = Ofs = Ocn = Ocs	0 dB
TTT	40 ms

*: For the experimentation in scenario 1 only. In scenario 2, we define a shadowing profile in the channel emulator to emulate UE mobility.

We have defined two scenarios. Scenario 1 uses a variable attenuator connected to the Tx output of the source eNB. In order to activate the HO condition, we manually increase the DL attenuation of the serving cell in 1 dB steps, which results in a decrease in the RSRP measured by the UE, as depicted in Figure 7. The lower peak preceding each plateau is due to the attenuation controller. Since the duration of the peaks is between 1 and $7ms$, this has no impact on the overall handover process. Indeed, the RSRP measurement and filtering do not take in account the previous measurements and the duration of the peak is less than TTT ($40ms$ in our setup).

Scenario 2 uses a Propsim F8 channel emulator in order to have a realistic channel, i.e., the 3GPP Extended Pedestrian A (EPA) channel model, with a 5Hz doppler. The corresponding RSRP measurement is depicted in Figure 8. Mobility is simulated by a shadowing profile that regulates the attenuation between all transceivers: the signals from the source eNB become weaker and the signals from the target eNB become stronger, which is a more realistic scenario than scenario 1 where the only signal whose strength changes over time is the Tx signal from the source eNB to the UE.

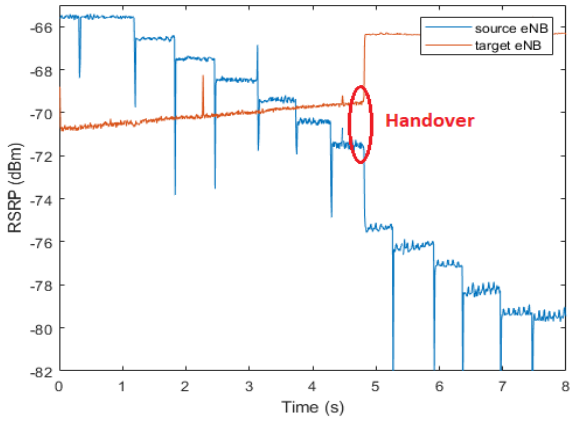

Fig. 7. RSRP measured over time, post L3 filtering, scenario 1

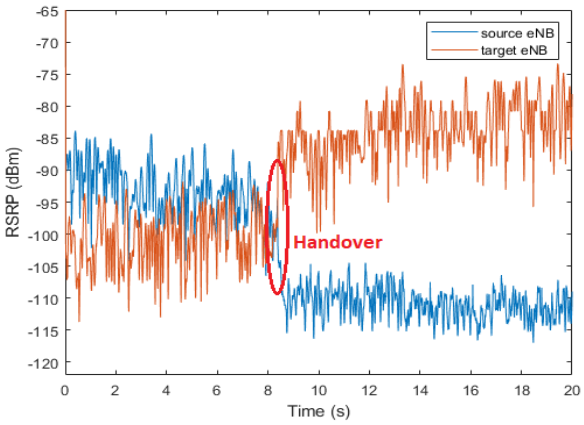

Fig. 8. RSRP measured over time, post L3 filtering, scenario 2

In Figure 7, at $t = 3.73s$, the RSRP of the target eNB is greater than the RSRP of the source eNB, but it is not high enough to enter the Equation 1 condition. At $t = 4.64s$, the HO condition is satisfied and there are no leaving condition (Equation 2) for $40ms$. Therefore, the UE sends a measurement report to the serving cell. At $t = 4.76$s, the UE receives the HO command, the HO procedure starts and the UE connects to the target eNB. While the UE is performing the HO, it does not measure the RSRP. Measurements after $t = 4.76s$ are made after the HO. Between $t = 4s$ and $t = 5s$, we observe a $3dB$ jump or drop in the measured RSRP of the target eNB and the source eNB respectively. This is due to the fact that the HO procedure induces a frequency change (i.e. an offset in the carrier frequency of the experimental setup, even through the DL frequency is supposed to be the same).

B. Validation of Handover Procedure

Figure 9 is a screenshot of a part of the OAI GUI with several lines. The top 3 lines represent the downlink control, its corresponding ACK and NACK respectively; and the bottom 2 lines represent the uplink control and its corresponding ACK respectively. The regions highlighted in red and yellow correspond to the periods when the UE is connected to the source and target eNB, respectively. The different steps of the HO procedure are: *1.* The UE launches its random access to the source eNB (RRC: idle to connected). *2.* The UE is connected to the source eNB without CN connection (RRC

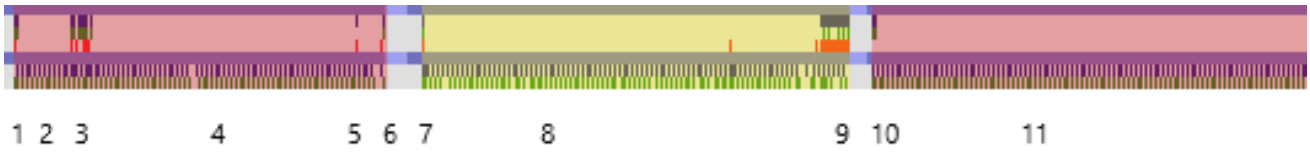

Fig. 9. OAI GUI capturing two HOs from the UE perspective

connected). *3.* After NAS exchanges, the UE is connected to the CN (RRC connected). *4.* The UE is connected to the source eNB with core connection (RRC connected) : Stable UL/DL. *5.* The UE sends a measurement report to the source eNB (RRC connected). *6.* The UE receives a *RRC connection reconfiguration* (RRC connected to idle). Start of HIT, start of HO. *7.* The *RRC connection reconfiguration* is completed to the target eNB (RRC idle to connected). End of HIT. *8.* The UE is connected to the target eNB (RRC connected). *9.* The UE sends a measurement report to the target eNB + retransmissions due to failures + reception of the *RRC connection reconfiguration*. *10.* The RRC connection reconfiguration is complete to the source eNB. End of the second HO. *11.* The UE is connected to the source eNB.

C. Performance Evaluation

Figure 10 shows the latency decomposition according to the different times defined in Equation 7. The duration of the three HO phases is $83ms$, $150ms$ and $8ms$ for preparation, execution and completion phases respectively. As stated in II-B4, the end of the execution phase and the beginning of the completion phase are done in parallel, resulting in a gain of $3ms$. Thus the total HO time is $238ms$.

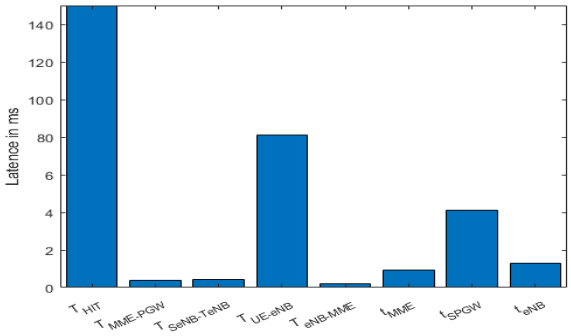

Fig. 10. Latencies measured

Specifically for the RRC connection reconfiguration time, T_{HIT} includes the duration of CFO estimation and the signal synchronisation (i.e. $124ms$) and the CFRA procedure (i.e. $12ms$). We use a single channel UE USRP board, so the CFO cannot be estimated on-the-fly. The UE scans CFO after breaking the connection to the source eNB, resulting in a reasonably large T_{HIT}. The value of $12ms$ depends on the PRACH configuration index chosen by the eNB. Indeed, after synchronisation, the UE waits for an RA procedure opportunity. We choose the PRACH configuration index 0 which offers an RA opportunity only at subframe 1 of even subframes, i.e. once every $20ms$. The UE synchronizes with the eNB by decoding the primary and secondary synchronization signals (PSS/SSS) that are broadcasted in subframe 0 and 5. Thus, if the UE synchronizes to an even-numbered frame at subframe 5, the opportunity to send a preamble will be at subframe 1 of frame N+2, $16ms$ later.

Moreover, in our experimental setup, there is no PCRF. Since the T_{IP-CAN} component consists of the processing time in the PCRF and P-GW and we know the exact processing time in the P-GW, which is $4ms$, assuming [2] measurements are consistent with our data, we can estimate the processing time in the PCRF as $20ms - 4ms = 16ms$.

Figure 11 shows the end-to-end uplink throughput, using the client-server tool *iperf* for generating UDP traffic. When connected to the CN, the UE receives an IPv4 address. To generate the uplink traffic, we bind the client to the IP address of the UE and the server to a machine routed to UE via CN. The throughput is measured on the server side. We first see that the throughput before handover is about 7Mbps. We can see a drop in throughput at $t = 2.6s$, a complete loss of traffic at $t = 2.7s$ and a start of recovery at $t = 2.8s$ to finally get back to the the cruising speed at $t = 2.9s$. The $0.2ms$ between the drop in throughput and the recovery corresponds to the duration of the handover.

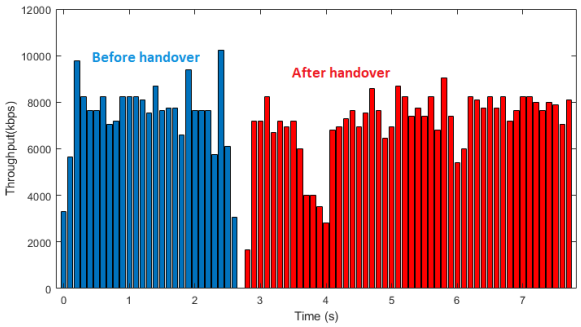

Fig. 11. End to End uplink throughput measurement

IV. Conclusion and Future Work

In this paper, we present an experimental testbed of the X2 handover procedure using an accessible and reconfigurable software defined radio environment with an end-to-end architecture (i.e. including both the radio access network and the core network). We find that most latencies are of the same order as those found in the state of the art of real experiments. The configurability and openness of the setup, however, brings a tradeoff, namely the latency caused by the reconfiguration of the UE to align its frequency with the target CFO.

Last but not least, OAI currently only supports Non-StandAlone (NSA) network with COTS UEs. In future work, we plan to integrate the X2 handover procedure into 5G to enable 5G NSA using an OAI UE instead of a COTS UE.

Acknowledgments

This work was partially supported by the ECSEL Joint Undertaking (JU) programme, under grant number N°826276 (CPS4EU project).

References

[1] M. Tayyab, X. Gelabert, and R. Jäntti, "A survey on handover management: From lte to nr," *IEEE Access*, vol. 7, pp. 118907–118930, 2019.

[2] D. Han, S. Shin, H. Cho, J.-M. Chung, D. Ok, and I. Hwang, "Measurement and stochastic modeling of handover delay and interruption time of smartphone real-time applications on lte networks," *IEEE communications magazine*, vol. 53, no. 3, pp. 173–181, 2015.

[3] B. Kang, N. Vijaykrishnan, M. J. Irwin, and T. Theocharides, "Power-efficient implementation of turbo decoder in sdr system," in *IEEE International SOC Conference, 2004. Proceedings.*, pp. 119–122, 2004.

[4] N. Nikaein, M. K. Marina, S. Manickam, A. Dawson, R. Knopp, and C. Bonnet, "Openairinterface: A flexible platform for 5g research," *ACM SIGCOMM Computer Communication Review*, pp. 33–38, 2014.

[5] K. Alexandris, N. Nikaein, R. Knopp, and C. Bonnet, "Analyzing x2 handover in lte/lte-a," in *2016 14th international symposium on modeling and optimization in mobile, ad hoc, and wireless networks (WiOpt)*, pp. 1–7, IEEE, 2016.

[6] J. Manco, G. G. Baños, J. Härri, and M. Sepulcre, "Prototyping v2x applications in large-scale scenarios using openairinterface," in *2020 IEEE Vehicular Networking Conference (VNC)*, pp. 1–4, 2020.