A Study on 5G Performance and Fast Conditional Handover for Public Transit Systems*

Claudio Fiandrino[a], David Juárez Martínez-Villanueva[a], Joerg Widmer[a]

[a] *IMDEA Networks Institute, Madrid, Spain*

Abstract

Fifth-generation (5G) networks are now in a stable phase in terms of commercial release. 5G design is flexible to support a diverse range of radio bands (i.e., low-, mid-, and high-band) and application requirements. Since its initial roll-out in 2019, extensive measurements studies have revealed key aspects of commercial 5G deployments (e.g., coverage, signal strength, throughput, latency, handover, and power consumption among others) for several scenarios (e.g., pedestrian and car mobility, mid-, and high-bands, etc.). In this paper, we make a twofold contribution. First, we carry out an in-depth measurement study of 5G in a large public bus transit system in a major European city. Second, based on the insights observed with the measurement study, we propose a new target cell selection criteria applicable to Fast Conditional Handover (FCHO), a 3GPP-specific 5G technique to foster reliable mobility. Our results are based on an extensive measurement campaign performed with several mobile phones connected to several mobile network operators totaling more than 1500 km over three months. The measurements reveal how flexible the network deployment is by analyzing Radio Resource Control (RRC) messages, mobility management and the suitability of our FCHO solution, and application performance.

Keywords: 5G, Cellular Network Measurement, Public Transit System

*The present submission is an extension of our paper entitled "Uncovering 5G Performance on Public Transit Systems with an App-based Measurement Study" that appeared in the Proceedings of ACM MSWiM 2022 [1].

Email addresses: claudio.fiandrino@imdea.org (Claudio Fiandrino), david.juarez@imdea.org (David Juárez Martínez-Villanueva), joerg.widmer@imdea.org (Joerg Widmer)

1. Introduction

The widespread commercial roll-out of Fifth-generation (5G) networks is nowadays a reality. 5G deployments can be either Non-Standalone (NSA) or Standalone (SA), depending on whether the legacy 4G or 5G infrastructure is used for control operations. At the radio access level, 3GPP has specified[1] that 5G New Radio (NR) can operate at different radio bands, i.e., Frequency Range 1 (FR1) which includes low-band (below 1 GHz) and mid-band (between 1 and 6 GHz) frequencies, and Frequency Range 2 (FR2) with high-bands at millimeter-wave frequencies (above 24 GHz). The design of 5G networks strives at providing flexibility to support highly diverse application requirements that stem from bandwidth-intensive applications like 4K/8K video, Augmented/Virtual Reality to massive machine-to-machine communication for (industrial) IoT, to ultra-reliable and low latency communication (URLLC) for teleoperation [2].

Background on 5G Measurement Studies. Since the initial 5G commercial roll-out in 2019, the research community has carried out extensive measurement studies *in the wild* to understand 5G operation [3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]. The literature on measurement studies carried out in Europe is thin [5, 12, 15]. By contrast, the vast majority of the literature [3, 4, 6, 7, 8, 9, 10, 11, 13, 14] focuses on 5G U.S. mid-band and high-band deployments and uncovered key aspects related to coverage, latency, throughput, and application performance. Some works aim at revealing network configuration parameters related to the management of FR2 deployments [10], understanding how predictable the throughput is at millimeter wave (mmWave) frequencies [6] and its implications to latency [13], mobility management [11] and to applications like video streaming [7, 8]. Other works dig into the dynamics of power management [4, 7] by breaking down the contribution of the radio at each state of the RRC state machine and of the mobility management by analyzing the

[1] Already in Release 15 (https://www.3gpp.org/release-15).

behavior of radio state transitions with different mobility patterns (stationary, walking, and driving) [7].

Motivations and Objectives of this Study. In this paper, we specifically study a unique feature of 5G mobility management, i.e., the conditional handover (CHO) [16] and its evolution Fast CHO (FCHO) [17]. CHO was first conceived in 3GPP Rel. 15 in 2017 [18] with the ultimate goal of improving reliability. The key intuition is to decouple two phases of the legacy handover procedure, i.e., *preparation* and *execution*: the preparation starts earlier than usual so that the instructions reach the users when are still in favorable radio conditions. Then, the actual handover is conditionally executed only if the conditions of the target cell become good enough. FCHO retains CHO candidates after the handover execution which enables the reuse of target cell preparation and reduces overheads. As the 3GPP specifications are not specific in regards to the criteria to choose the target cell upon meeting the conditions, in this work we explore how to make this process efficient based on the specific characteristics of mobile network connectivity on bus-based mobility. For this, we use hypergraphs to study the cell attachment problem over time and propose an effective target cell selection criteria.

We ground our study on CHO and FCHO on the insights obtained with our extensive 5G measurement study for a major European city, Madrid [1]. In such a city, the existing mobile network operators have rolled out 5G NSA over mid-bands. Unlike the above works on 5G measurement studies, we have studied 5G performance in the public bus transit system and we specifically analyzed the unique features of 5G mobility management. The reason to study the public bus transit system is twofold. *First*, despite a reduction of use because of the COVID-19 pandemic [19], buses are still a widely used form of public transport in Europe and Canada, more than in the U.S. [20]. The existing measurement studies can not be applied to infer 5G performance on the bus public transit system because the transportation modes [21] and travel times [22] of buses are radically different from those of cars studied so far. Other 5G studies investigate low-bands in rural environments [23], which is also orthogonal to our work.

Second, 5G-fueled teleoperated vehicles are regarded as important in future city-wide transport systems. To extend the limited use of teleoperated buses within factories and campuses [24], a measurement study like ours is essential to uncover problems and challenges that are hard to study through simulations or testbeds due to the complexity of the scenario. All the above arguments make our analysis fully orthogonal to previous measurement studies. Regarding the analysis of mobile network performance over public bus transit systems, unlike ours, the study of Elsherbiny et al. focuses on LTE [25]. Finally, Pan et al.[26] focus on 5G performance on high speed trains.

Key Contributions and Findings. Our objective with this paper is to contribute toward the understanding of real-world 5G deployments and their performance when using the public bus transit system in a major European city. The key contributions and findings of our study are summarized as follows:

C1. We define and use an app-based methodology to collect a rich set of data in a major city-wide scenario. We present a thorough analysis that allows an understanding of the configuration of the network deployment, coverage, mobility, and end-to-end application performance of several mobile network operators.

C2. We study the specific case of the public bus transit system the mobile network performance in bus routes traversing the Madrid city center, the suburban municipalities nearby Madrid, and the interconnections between the latter with the city center.

C3. We study specifically the mobility management for the case of public bus transit and propose a new target cell selection criteria for the case of FCHO, a distinct feature of 5G mobility management.

C4. We will release the artifacts of the current study in the existing repository containing measurements and code of the preliminary conference version [1]: `https://git2.networks.imdea.org/wng/5g-bus-public-transit`.

F1. We find that the mobile network operators keep the same configuration

across space (urban and suburban areas) and time and that the flexibility of 5G NR is yet to be unleashed. For example, when using applications with different bandwidth requirements, we observe no changes in numerology configuration.

F2. We find that across different days in a week, users attach frequently to the same set of Base Stations (BSs) during bus routes, which makes bus mobility amenable to CHO, a recently introduced feature in 3GPP Release 16.

F3. We find that the 5G deployments vary significantly among the studied operators and this translates into different end-to-end user performances when riding a bus on an urban or a suburban route.

Paper Organization. The rest of the paper is organized as follows. Section 2 outlines the data collection methodology, which includes the descriptions of the dataset and the applications utilized. Section 3 delves into the 5G network parameter configurations that the monitored mobile network operators use for their own network. Section 4 explores mobile network deployments with a focus on urban and suburban zones of the metropolitan area of Madrid and provides a preliminary analysis of mobility management. Based on such analysis, Section 5 digs deep into the suitability of CHO and FCHO for the case of a public transport system and proposes an efficient target cell selection strategy. Section 6 investigates end-to-end network performance and finally, Section 7 provides concluding remarks.

Ethical Considerations. This study was carried out by paid and volunteer personnel. No personally identifiable information (PII) was collected or used, nor were any human subjects involved. We purchased multiple cellular data plans from major EU mobile network operators. Our study complies with the wireless carriers' customer agreements. This work does not raise any ethical issues.

2. Data Collection Methodology

The 5G ecosystem is constantly expanding and evolving. Since its first commercial roll-out in 2019, coverage, and 5G devices have now become mature. **Measurement Tools**. We use multiple smartphone models with 5G support and diverse specifications: Xiaomi Mi Mix 3 5G (M1810E5GG), Xiaomi Mi 10 (M2001J2G), and Samsung Galaxy S20 Ultra 5G (SM-G988U). Compared to the M1810E5GG, the M2001J2G and SM-G988U have a superior 5G modem and chipset, increased CPU frequencies, and more RAM.

For our measurements, we use several Android applications: each of them supplies different information. Before taking actual measurements, we performed extensive tests to understand the potential and limitations of joint use of the apps. Some of the apps require root access, hence we rooted all the phones.

• ***MobileInsight*** [27] exposes over-the-air messages from the chipset to user space by exploiting the so-called diagnostic mode, a second channel between the hardware chipset and software. This allows collecting control- and user-plane protocol interactions that reveals parameter configurations with which the 5G NR BSs[2] instructs the User Equipments (UEs). The current support for 5G is limited to RRC messages. We use this application to reveal the configurations of the mobile network operators in § 3.

• ***5G tracker*** [28] is developed by the University of Minnesota with a free license and allows to record active and passive 5G measurements, including radio type (5G/4G/3G), Cell ID, Physical Cell ID (PCI), NR Bands, ARFCN, Tracking Area Code (TAC), Mobile Network Code (MNC), Total RX/TX Bytes, Signal Strength Info (RSSI, RSRP, SINR, RSRQ), movement speed, throughput and latency with built-in iperf and ping tools.

• ***GNetTrack*** [29] is developed by Gyokov Solutions and provides similar parameters to 5G Tracker. Unlike 5G Tracker, GNetTrack is more stable during

[2]In this paper we use the term BS to identify both LTE (evolved Node B (eNB)) or 5G NR (next Generation Node B (gNB)) stations when the Radio Access Technology (RAT) is not key for the discussion.

long recordings. The pro version comes with a small one-off fee but is instrumental to download the logged data. We use this application for the study of coverage and mobility management in § 4.

• ***Network Signal Guru (NSG)*** [30] is developed by Qtrun Technologies and provides extensive information. In addition to the parameters listed for 5G Tracker, NSG also reports MIMO configuration, beam index, block error rate (BLER), modulation and coding scheme, and the number of allocated RBs among others. Unfortunately, it comes with a monthly fee and it is extremely difficult to export the logged data without proprietary expensive software.

• ***Ookla's Speedtest*** [31] is the state-of-the-art tool for assessing connection status over the Internet and allows benchmarking network throughput, latency, and video streaming quality. We use this application for the end-to-end network performance tests in § 6.

Mobile Network Operators and Methodology. We select three major EU mobile network operators that have deployed 5G in several bands like mid (3.3-3.8 GHz, band n78) and low bands (700 MHz, n28). High (mmWave) bands (24.25-27.5 GHz) are not deployed yet. The measurement study is conducted in areas where only mid-bands are available.

We collect both 4G and 5G measurements over three months (Feb. - May 2022) in both morning and afternoon in Madrid and its suburban municipalities with the following two main objectives: to *i)* understand 5G performance in a large public bus transit system *ii)* compare the deployments and performance of several operators.

Our dataset has been collected with field trips totaling 1677 km and having observed a total of 4167 unique cells including both primary and neighbor cells (specifically, the numbers of unique primary cells per operator are respectively 1624, 1553, and 990 for Op. 1, Op. 2 and Op. 3). We defined three main trajectories (Fig. 1 outlines all of them together). The first one consists of 8 bus routes (4 suburban and 4 urban[3]) that start from a bus stop nearby the IMDEA

[3]We define as urban those bus routes that travel with the diameter of the orbital highway

Figure 1: Map of Madrid with bus routes

Networks facility in Leganés (a suburban municipality south of Madrid), reach the city center through a highway, navigate Madrid's touristic and business areas and come back. In addition, we use two circular bus routes for the city center and one that interconnects several municipalities in the south of Madrid.

3. Mobile Network Configuration (under Mobility)

The configuration of BSs is essential to optimize the Quality of Experience from the user perspective and maximize the revenue of the operator. This task is extremely complex as BSs handle mobility and handovers, load balancing, and interference among other tasks. Overall, there exist on the order of thousands of configuration parameters for an LTE eNB [32]. 5G NR retains the basic LTE frame structure (10 ms frame and 1 ms subframe), but augments flexibility in the radio, e.g., with different sub-carrier spacing or numerologies among numerous other parameters. Indeed, to support different operating frequencies (i.e., low-band, mid-band, and high-band), it is extremely hard to use a unique sub-carrier spacing without sacrificing too much efficiency or performance.

In this section, we answer to the following questions:

M-30 circling all the central districts of Madrid municipality.

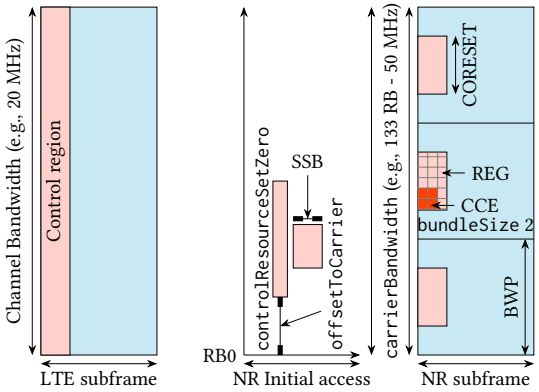

Figure 2: LTE and 5G NR parameter configuration

- *How much of such radio flexibility is available today?*
- *Do operators share the same network configuration?*

To answer these questions, we exploit MobileInsight [27] version 6.0.0 which allows analyzing 5G RRC operation. The collection and analysis of Over-The-Air (OTA) packets reveal the configuration of key parameters. Fig. 2 shows the difference between LTE and 5G NR at the radio level configuration and be utilized to guide the reader through this section. In a nutshell, the main difference lies in the fact that 5G NR enables flexibility by dividing into parts the bandwidth of both control and data.

Initial Access. The details of the full initial access procedure are well-described in [33]. For the scope of this paper, we summarize the procedure and analyze the configuration parameters of the various carriers. During the initial access, a UE acquires from the Synchronization Signal Block (SSB) the Master Information Block (MIB) which provides the UE with synchronization and information on frequency and time resources required to derive the System Information Block 1 (SIB1). SIB1 information is crucial for initial access as it contains random access configuration and thresholds for minimum measured channel quality and received power. The aforementioned frequency and time resources are respectively termed `controlResourceSetZero` and `searchSpaceZero`. The former parameter allows to lookup in 3GPP TS38.213, Table 13-4 [34] the amount

Table 1: Initial access parameter configuration

Parameter	Operator		
	Op. 1	Op. 2	Op. 3
`controlResourceSetZero`	11	12	10
`searchSpaceZero`	2	2	4
`offsetToCarrier`	14	16	12

of Radio Blocks (RB) allocated for the MIB and the offset in the number of RBs from the `offsetToCarrier`. `searchSpaceZero` instead allows to look up in 3GPP TS38.213, Table 13-11 [34] the parameters required to determine the starting slot and system frame number (SFN) essential to operate in the network.

Table 1 shows how operators configure the parameters to decode the MIB. The results are computed over all the monitored BSs and duration of the data collection campaign. The operators configure statically these parameters across space and time. The configuration of `controlResourceSetZero` leads to the same frequency allocation (48 contiguous RBs and 1 OFDM symbol for the PDCCH) and the only difference is the offset from the RB0 (i.e., 12, 14, and 16 RBs for Op. 3, Op. 1 and Op. 2 respectively). Regarding the time, the only difference between the operators is the slot offset relative to the start of the frame from which to start monitoring MIB1.

Numerology and Bandwidth Parts. In 5G NR, the maximum bandwidth available in FR1 and FR2 is respectively 100 MHz and 400 MHz. This allows the allocation of contiguous blocks of the spectrum which is beneficial for bandwidth-intensive applications. However, scanning continuously such bandwidth would be power-costly for modems, especially when serving non-bandwidth intensive applications. Therefore 5G NR features *bandwidth parts* (BWP), i.e., smaller portions of contiguous RB of the entire bandwidth (`carrierBandwidth`) that are assigned to the UEs via RRC signaling. Each BWP is configured with a specific setting like sub-carrier spacing (`subcarrierSpacing`) and location (`locationAndBandwidth`). A UE can only access BWPs that are assigned to it

Table 2: Bandwidth part configuration

Application	Parameter	Operator		
		Op. 1	Op. 2	Op. 3
Ping	`subcarrierSpacing`	30 kHz	30 kHz	30 kHz
	`carrierBandwidth`	133	245	162
	`locationAndBandwidth`	36300	[8799, 12943]	31624
Iperf	`subcarrierSpacing`	30 kHz	30 kHz	30 kHz
	`carrierBandwidth`	133	245	162
	`locationAndBandwidth`	36300	[8799, 12943]	31624
SpeedTest	`subcarrierSpacing`	30 kHz	30 kHz	30 kHz
	`carrierBandwidth`	133	245	162
	`locationAndBandwidth`	36300	[8799, 12943]	31624

(up to 4 in uplink and 4 in downlink).

Table 2 shows the list of configuration parameters above mentioned for the different operators when using different applications. From Table 2 we draw two conclusions. First, different applications (ping generates light traffic while iperf was configured to carry a significant amount of traffic) are not treated differently by the network. Second, similarly to the initial access parameters, numerology, and bandwidth parts are kept with the same configuration across network space (i.e., BSs) and time. Specifically, a subcarrier spacing $\Delta f = 30$ kHz indicates a numerology of $\mu = 1$ because: $\Delta f = 2^{\mu} \cdot 15$ kHz (by contrast LTE only features one subcarrier spacing of 15 kHz). We observe that the operators use different `carrierBandwidth`s. As the `subcarrierSpacing` is 30 kHz, the carrier bandwidths are 50 MHz (133 resource blocks), 90 MHz (245 resource blocks), and 60 MHz (162 resource blocks) for Op. 1, Op. 2, and Op. 3 respectively.

Control Resource Set. Each BWP specifies its own control region where the mobile device searches downlink control signals. 5G NR Control Resource Set (CORESET) differs from that of LTE in a number of ways. First, the resources can be allocated both in time (number of OFDM symbols) and frequency (part

Table 3: Coreset configuration

Parameter	Operator		
	Op. 1	Op. 2	Op. 3
`controlResourceSetId`	[1, 1, 1]	[0, 0, 0, 0, 1, 1]	[0, 1, 1]
`bundlesize`	n4 (0)	n4 (0)	n4 (0)
`cce-REG-MappingType`	nonInterleaved (1)	nonInterleaved (1)	nonInterleaved (1)
`searchSpaceId`	[1, 2]	[1 8 4 9 5]	[1, 2]
`monitoringSlotPeriodicity AndOffset`	[sl1 (0), sl1 (0)]	[sl1 (0), sl40 (8), sl1 (0), sl40 (8), sl1 (0)]	[sl1 (0), sl1 (0)]

of the channel bandwidth) while in LTE the resources are only configured in time because the control region occupies the whole channel bandwidth. Second, 5G NR features two types of CORESET: common and UE-specific. A Common Search Space (CSS) is shared across all the UEs (i.e., `controlResourceSetId` equal to 0) while a UE-specific search space (USS) is configured per UE basis (i.e., `controlResourceSetId` different than 0). The smallest resource unit for the CORESET allocation is the Resource Element (RE) which consists of one subcarrier in frequency and one OFDM symbol in time. A Resource Element Group (REG) consists of 12 RE, and a REG bundle consists of one or more REGs defined by the parameter `bundleSize`. The `searchSpaceId` defines how many candidate locations the UE can perform decoding at each bundle. The number of `searchSpaceId`s is one less than the `controlResourceSetId`. The Control Channel Element (CCE) is a combination of multiple REGs and the number of REG bundles can vary within one CCE. This relationship is defined by the parameter `cce-REG-MappingType` and can be either interleaved or not interleaved. For the latter type of mapping, all CCEs for a DCI with a given `bundleSize` are mapped in consecutive REG bundles of the CORESET. Conversely, an interleaved `cce-REG-MappingType` can enable both a time-domain processing gain and frequency-domain diversity. `monitoringSlotPeriodicityAndOffset` defines how often the UE performs monitoring.

Table 3 shows Op. 1 and Op. 3 share a similar configuration with the only exception of `controlResourceSetId` where Op. 1 has only configured UE-specific

space searches. Op. 2 instead, given also the larger bandwidth at disposal (see Tab. 2) features a diverse configuration with multiple common search spaces and locations where UEs can perform decoding of the downlink control channel. Note that the different spaces are monitored more or less often: sl1 (0) indicates a per slot monitoring, sl40 (8) indicates that the UE monitors the search space every 8 slots.

Handover. A distinct feature of mobile networks is the way in which they maintain connectivity during mobility. Traditional approaches perform handover from one BS to another based on measurements of radio quality which may include criteria like Radio Signal Strength Indicator (RSSI), Signal to Interference and Noise Ratio (SINR), Reference Signal Received Quality (RSRQ), Reference Signal Received Power (RSRP) [11]. Specifically, a serving or source BS informs the devices about when and how often measurements of the radio quality should be reported. Some of the key configuration parameters for the process are `reportInterval`, `MaxReportCells`, `timeToTrigger` and `reportAmount`. These specify respectively the time between periodic measurements, the maximum number of non-serving cells to be included in the report, the time during which specific criteria need to be met to trigger a measurement report and the number of measurement reports applicable for both event-based and periodic reports. For the latter parameter, all the operators configure `reportAmount` as `infinity (7)` (periodic) or `r1 (0)` (event-based). Table 4 shows the details of the configuration parameters, including relevant timers like `t304` and `t310` that deal with mobility and upon expiration trigger connection re-establishment procedures. Specifically, `t304` starts upon reception of a reconfiguration message and stops upon successful configuration at RRC level. Instead, `t310` starts upon detection of physical layer problems at the serving cells and stops upon meeting different conditions like receiving a predefined number of in-sync messages from the serving cell or initiating a reconfiguration procedure. We can observe from Table 4 that the configurations vary quite a lot among the operators, including the timers.

Table 4: Handover configuration. N/A denotes fields of Op. 1 for which we could not decode any information.

Parameter	Operator		
	Op. 1	Op. 2	Op. 3
`reportInterval`	N/A	ms120 (0), ms480 (2)	ms240 (1), ms5120 (6)
`MaxReportCells`	N/A	1	1, 4, 8
`timeToTrigger`	N/A	ms160 (6), ms640 (11)	ms160 (6), ms320 (8), ms640 (11)
`t304`	ms1000 (5)	ms1000 (5)	ms500 (4)
`t310`	ms2000 (6)	ms2000 (6)	ms1000 (5)

4. Mobile Network Deployment and Mobility Management

While the extensive measurement studies *in the wild* have shown the details of 5G for the U.S. [3, 4, 6, 7, 8, 9, 10, 13], to the best of our knowledge there are no existing studies of 5G performance on public transit systems for European cities. Our analysis is limited by the fact that to the current date, the mobile network operators analyzed have not deployed FR2 or SA deployments. In this section, we, therefore, answer the following questions:

- *How is the mobile network deployment in both urban and suburban environments that are served by the public bus transit system?*
- *How is the specific bus-based mobility managed by the mobile network operators?*

To answer these questions, we exploit GNetTrack [29] and try to characterize different aspects that affect mobility like network deployment and handover management for both urban and suburban scenarios.

4.1. Network Deployment

Connectivity. Because of pre-determined routes, we expect that moving on buses the mobile device attaches preferentially to the same BS. We investigate whether this is true with wheel charts where each semi-circle represents a different day (the most inner-circles are Mondays, the most outer-circles are Fridays) and

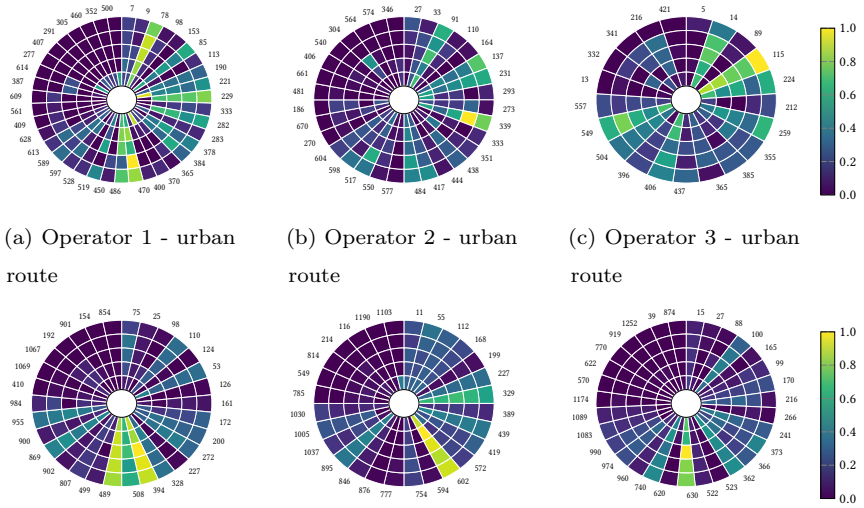

(a) Operator 1 - urban route

(b) Operator 2 - urban route

(c) Operator 3 - urban route

(d) Operator 1 - suburban route

(e) Operator 2 - suburban route

(f) Operator 3 - suburban route

Figure 3: BSs association over different days of a week

we depict clockwise how frequently the mobile phone associates to subsequently encountered BSs. As an urban scenario, we choose a route that follows an avenue traversing Madrid south-north. As a suburban scenario, we choose a route that interconnects the municipality of IMDEA Networks premises with a neighbor municipality, both outside the Madrid city center. For each trip, we count the time that the mobile device was connected to a given BS and then normalize with respect to the local maximum of each route. Values close to 1 indicate that the mobile was attached for a significant amount of time. Viceversa, values close to 0 indicate intermittent association. Fig. 3 shows the plots for each operator for the two routes. We can observe that the set of BSs can be categorized between *frequently* and *infrequently* seen BSs. For example, the BS with anonymized ID 78 in Fig 3(a) belongs to the former category while the BS with anonymized ID 854 in Fig 3(d) belongs to the latter category. The status of frequently or infrequently seen BSs holds across days. Under such predictable network status and bus mobility [35], conditional handover, a recently introduced 5G NR optimization, seems to find room for applicability [36].

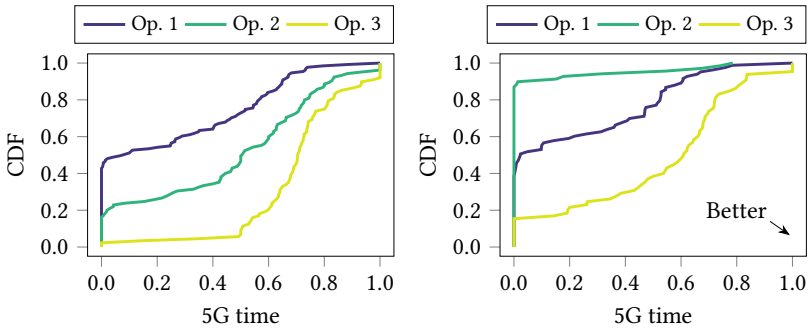

(a) Urban (b) Suburban

Figure 4: 5G time observed in different scenarios

With conditional handover (CHO), the handover preparation and execution are decoupled and the network prepares target cells in advance than with usual handover. Upon satisfying a condition criterion, the UE and not the network executes the handover. With Fast-CHO [17] the UE is allowed to re-use early prepared target cells instead of releasing them, which is suitable for bus mobility because of the following reasons. Bus mobility features predictable routes and schedules and this makes the UEs' mobility predictable. Frequent passengers may considerably benefit from FCHO given that they are going to re-use multiple times an already prepared target cell.

5G Time. Although 5G deployments are becoming more and more widespread to serve nearly 700 million subscribers, LTE is still the dominant mobile technology with 4.7 billion subscribers as of Q4 2021 [37]. To maintain compatibility, specific procedures allow to transfer of the ongoing connection in both directions, i.e., 4G $\rightarrow$ 5G [38] and 5G $\rightarrow$ 4G [39]. This is called inter-RAT, or vertical handover to differentiate it from intra-RAT or horizontal handover where the ongoing connection is transferred between two $4G$ or two $5G$ BSs. We find that Inter-RAT handovers typically occur on the same BS because operators typically deploy 5G NR antennas co-located with existing LTE infrastructure and align cell boundaries so that the PCI coincides [40].

Considering inter-RAT handovers on the same BS, we now study how long a connection sticks to 5G (termed t_{5G}) with respect to the total time of association

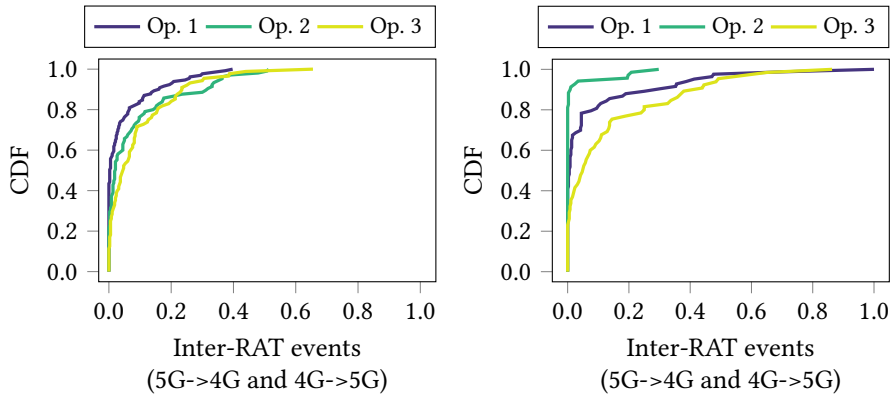

(a) Urban

(b) Suburban

Figure 5: Inter-RAT events

with the BS (termed $t_{4G} + t_{5G}$) with both RATs. Thus:

$$\text{5G Time} = \frac{t_{5G}}{t_{4G} + t_{5G}}. \tag{1}$$

Fig. 4 shows the result for the entire set of BSs (see § 2) in urban and suburban routes. Values of 5G Time close to 1 indicate that the user is served by a given BS mostly with 5G NR. Conversely, values of 5G Time close to 0 show that LTE is mainly used. With Op. 3, users benefit from 5G connectivity more often than any other operator and this statement holds for both urban and suburban scenarios. In the latter case, 17% of the BSs feature LTE connectivity only. Op. 2's 5G deployment is practically nonexistent in suburban areas with only a tiny fraction of BS offering 5G service (16%) while Op. 1 has deployed 5G nearly equally across the analyzed areas.

4.2. Handover Management

Inter-RAT and Intra-RAT events. We now study the distribution of inter-RAT events that occur on a single BS for each of the operators. While the 5G Time metric studies how long the connection uses the same RAT, such Inter-RAT events identify how frequently the RAT changes are. These can be attributed to B1 and B2 type of handover events [41]. We count the number of Inter-RAT events for each BS of urban and suburban routes and then normalize with respect

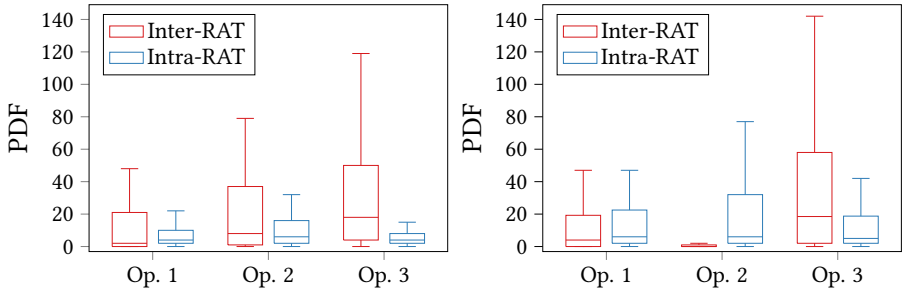

(a) Urban (b) Suburban

Figure 6: Inter- and Intra-RAT events

to the maximum. This allows providing a direct comparison between the two scenarios shown in Fig. 5 and Fig. 6. We can observe that there is a higher number of events in suburban areas than in urban areas, with the exception of Op. 2 because of its limited 5G deployment with respect to the other operators. Overall, in the suburban areas, the presence of 5G is lower and is more probable that the connection frequently moves across RATs.

Next, we generalize the analysis and consider both Intra-RAT and Inter-RAT events, and, for the latter category, we consider those occurring in the same and between different BS. Intra-RAT events are attributed to A1-A6 type of handover events [41]. Fig. 6 shows that the operators have different handover management. With Op. 3 handovers are more frequent than for the other operators and Inter-RAT events are much more common than Intra-RAT events regardless of the scenario. In comparison with Op. 3, handovers in Op. 1 are less frequent and Intra-RAT events are more common in both urban and suburban scenarios. Unlike the other operators, Op. 2 enforces different policies in the two scenarios: in urban routes, Inter-RAT events are more frequent, and, viceversa, in suburban routes Intra-RAT events are more frequent.

Ping-pong Events. Fig. 7 show the details of BS association and RAT usage over time (the actual BS PCIs are anonymized). As an urban scenario, we choose again for the analysis of BS connectivity the route following one of the main avenues. Likewise, we choose as a suburban scenario the route that interconnects

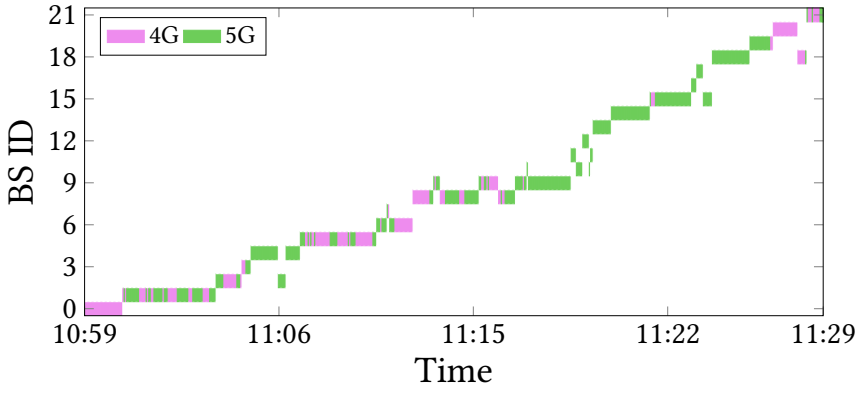

(a) Operator 2 in urban scenario

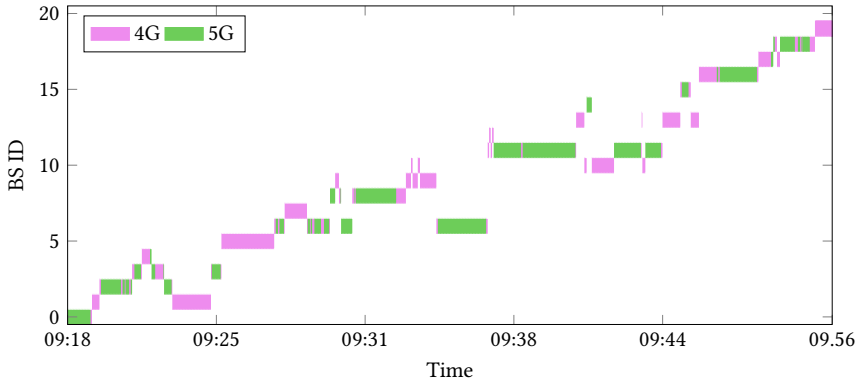

(b) Operator 1 in suburban scenario

Figure 7: BS association and type of RAT connectivity during two bus rides

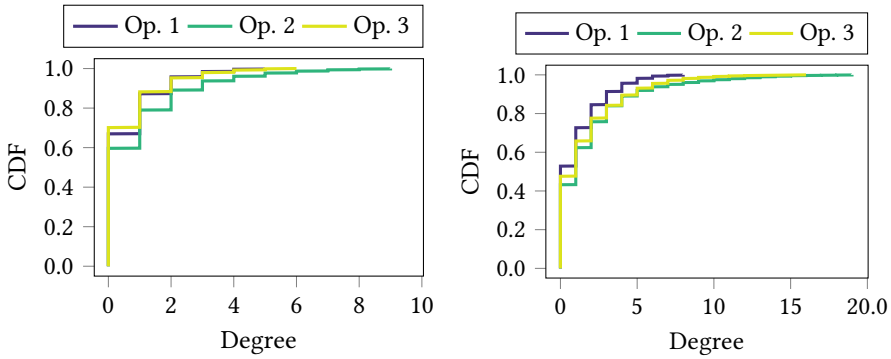

(a) Urban

(b) Suburban

Figure 8: Degree of ping-pong events

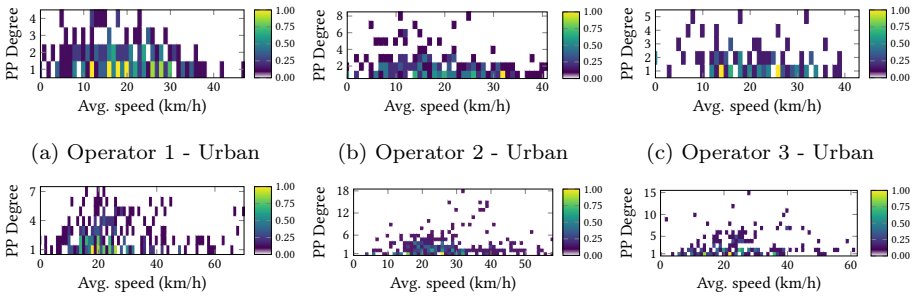

(a) Operator 1 - Urban (b) Operator 2 - Urban (c) Operator 3 - Urban

(d) Operator 1 - Suburban (e) Operator 2 - Suburban (f) Operator 3 - Suburban

Figure 9: Analysis of ping-pong events with respect to speed

the municipality of IMDEA Networks premises with a neighbor municipality still outside Madrid. In both cases, mainly in the suburban route, we can identify ping-pong effects, i.e., the connectivity is transferred to previously seen cells back and forth. For example, in Fig. 7(b), the anonymized BS with ID 6 is seen four times between 09:25 and 09:31. The underlying motivations for this behavior are two. *First*, the specific routes of buses might follow paths that make the bus traverse the same area at different times. In this case, the device camps on that cell for some time. The *second* reason might be a network sub-optimal handover decision and this typically happens for a very short time (see BS ID 10 in Fig. 7(b)).

We decided to take a closer look at ping-pong events. Specifically, we define the Ping-Pong (PP) degree as the hop number between the first and each subsequent connection to the same BS. For example, in Fig. 7(b), BS ID 2 has a PP degree of 6. We count all the hops and not just the maximum to fully characterize the PP behavior. A degree equal to 0 indicates no PP. Fig. 8 shows that *i)* PPs happen at least 40% and 60% of the time in urban and suburban routes respectively, and *ii)* Op. 1 is the operator with the lowest number of PP in both scenarios. Viceversa, with Op. 2 the PP degree reaches the highest values.

Next, in Fig. 9 we further dig deep into the problem. We compute the average speed of the mobile during the maximum PP degree time window and relate this to the PP degree metric. We find that the highest PP degrees (> 8) always occur in suburban areas for an average speed below 40 km/h. This occurs when

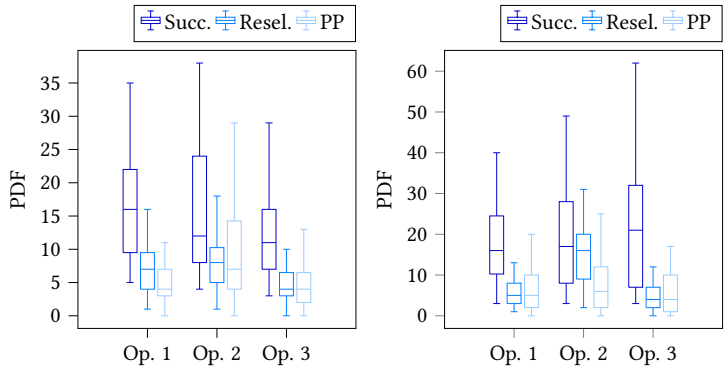

(a) Urban (b) Suburban

Figure 10: Degree of ping-pong events

the buses are driving within the suburban municipalities and not on the highway that interconnects those municipalities with Madrid. We attribute those events to the specific bus routes navigating in a same area. Regardless of the scenario, when the bus moves very slowly (< 10 km/h) or very fast (> 40 km/h) on the highway, PPs of any degree are rare which indicates that in such circumstances there might be sub-optimal network decisions.

Overall Event Analysis. Having analyzed Inter- and Intra-RAT and PP events, we now delve into the distribution of events that can be deemed as successful, those that lead to a PP with a degree equal to 2 and those events that are unsuccessful and lead to cell re-selection to the original BS (in other words, PP degree equal to one). Fig. 10 portrays the distribution of all the events across all measurement days grouped per urban and suburban routes. The results show that the distribution of successful events is higher in urban scenarios for Operator 1 and Operator 2 while for Operator 3 holds the contrary. Cell re-selection and PP account for a non-marginal amount of events, with PP usually higher in urban than in suburban scenarios. This is in line with the previous result shown in Fig. 8.

Reference Signal Received Power (RSRP). Fig. 11 shows that there is no difference in RSRP change among the operators in both urban and suburban scenarios. We measure the received signal power of the Synchronization Signal

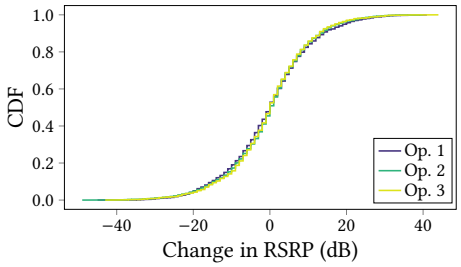

Figure 11: Change in RSPR before and after a handover event

(SS) and the RSRP is defined as the linear average over the power contribution of all the RE carrying (SS). We compute the RSRP change as the difference between the sample measures after and before the handover and the median is exactly 0. This means that it is equally probable to get an improvement or not.

We now dig deep into this result and explore changes in RSRP before and after a handover for different types of handover (i.e., inter- and intra-RAT), and according to the speed. Fig. 12, Fig. 13, and Fig. 14 categorize the results for urban and suburban cases for Operator 1, Operator 2, and Operator 3 respectively. Across all the scenarios, we can highlight the following general trends: *(i)* the RSRP usually improves for Inter-RAT 5G-4G events (the corresponding horizontal bars of the scatterplots are dominated by blue colors), *(ii)* the RSRP usually worsen for Inter-RAT 4G-5G events (the corresponding horizontal bars of the scatterplots are dominated by blue colors), *(iii)* the RSRP usually improves for 4G-4G Intra-RAT events and worsen for 5G-5G events, and *(iv)* there is no specific trend in regards to speed.

5. The Applicability of (Fast) Conditional Handover

5.1. A Primer on CHO and FCHO

CHO. CHO was firstly discussed during 3GPP Rel. 15 in 2017 [18] introduced as a part of 3GPP Release 16 (finalized in 2020 [42]) with the overarching goal of improving the reliability of legacy LTE handover procedure. While the latter consists of *preparation*, *execution*, and *completion*, CHO decouples *preparation* and *execution*. In such a way, the network achieves higher mobility robustness

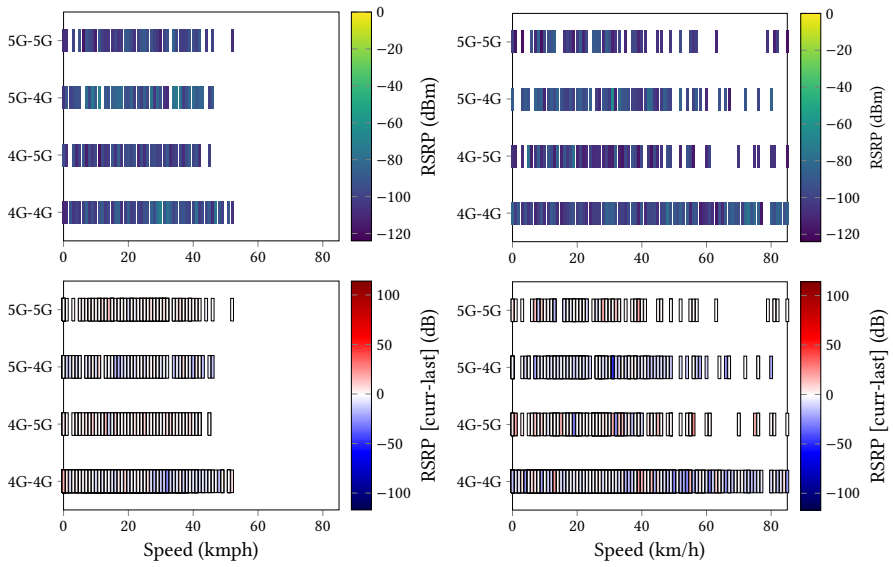

(a) Urban (b) Suburban

Figure 12: Breakdown of RSRP changes according to speed and different types of handover events (before - last - and after - curr -) for Operator 1

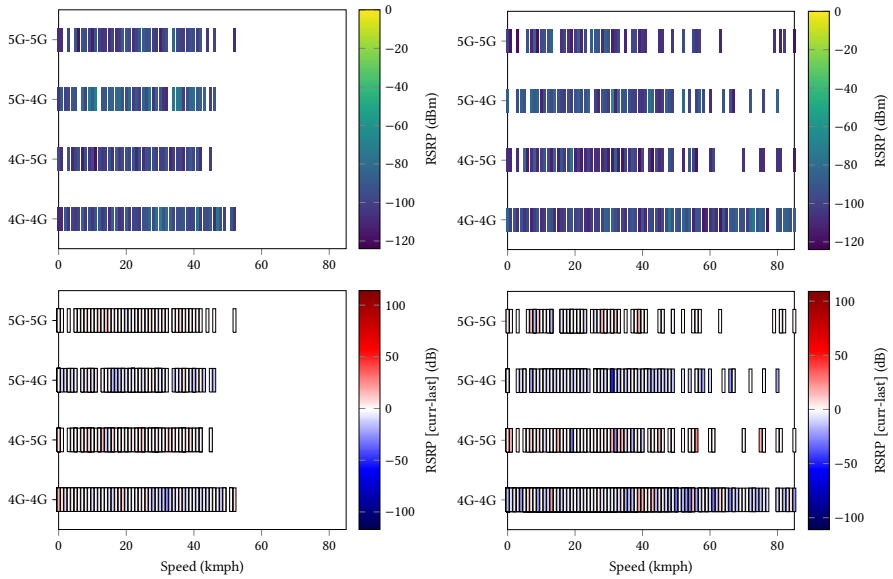

(a) Urban (b) Suburban

Figure 13: Breakdown of RSRP changes according to speed and different types of handover events (before - last - and after - curr -) for Operator 2

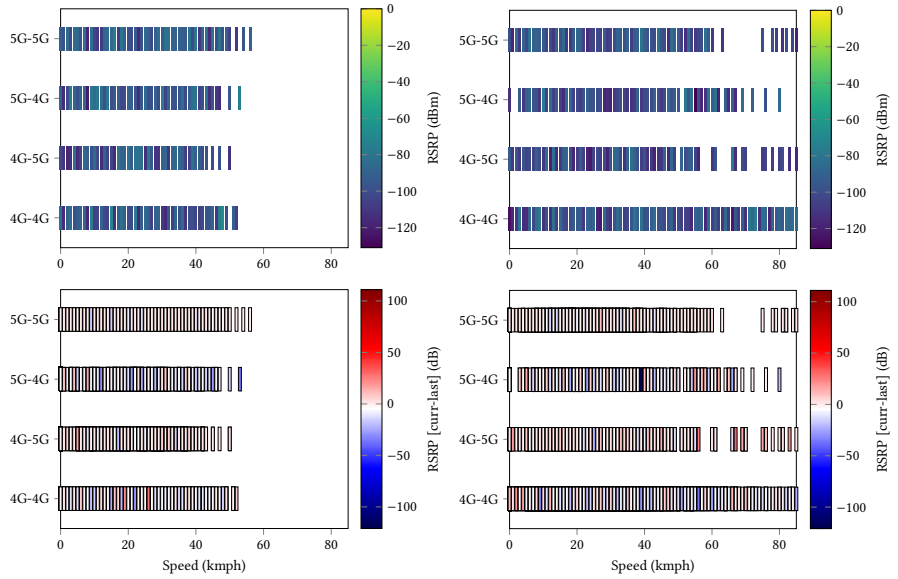

(a) Urban (b) Suburban

Figure 14: Breakdown of RSRP changes according to speed and different types of handover events (before - last - and after - curr -) for Operator 3

because the handover is prepared in advance so that the instructions reach the UE when is still in favorable radio conditions and the actual handover is conditionally executed only if the conditions of the target cell become good enough. Until now, CHO has been proven successful in several scenarios that include non-terrestrial networks [43], for integrated access and backhaul [44] and in 5G NR unlicensed [42].

FCHO. A fundamental drawback of CHO is that it introduces excessive signaling overhead [36], which is particularly a challenge in FR2 scenarios with dense cell deployments. 3GPP mandates that there may be up to eight candidate target cell prepared and the UE is configured with up to two conditions for the same reference signal. Despite this restricted action space, after each handover, the legacy CHO procedure releases the configurations for CHO preparation. By contrast, Fast CHO (FCHO) [45] maintains CHO candidates after the handover which enables the reuse of target cell preparations [17]. FCHO reduces *i)* failures in specific conditions like cell boundaries where the sharp increase of inter-cell

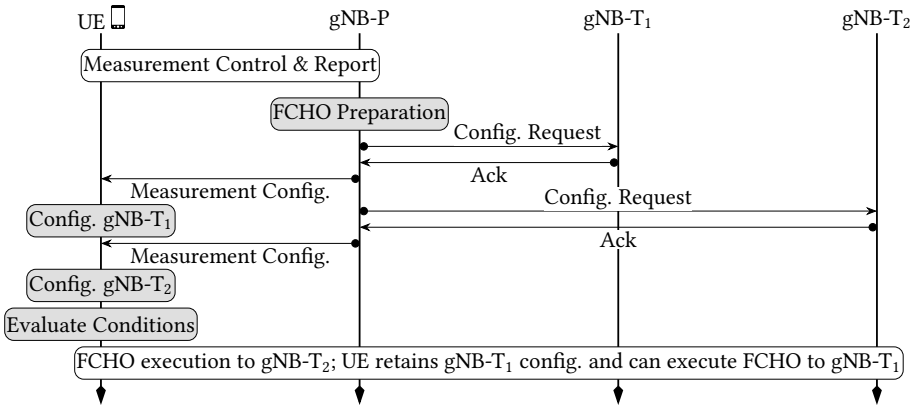

Figure 15: FCHO workflow with preparation to two gNBs

interference may negatively impact the capability of the network in receiving measurements and *ii)* signaling overhead because the target cell preparation is not always repeated after each successful handover. With the objective of reducing signaling other works [16, 46, 47] have exploited Deep Learning (DL) solutions to narrow down the number of candidate cells to prepare the handover to and select the most appropriate among the eight possible targets.

Next, we dive deep into the logic of FCHO, and then we study its suitability for the public transport system.

FCHO Deep Dive. Fig. 15 shows the workflow execution of FCHO for a UE currently served by a primary gNB (gNB-P in the Figure). The UE performs legacy measurement reports including potential candidate target cells and, upon meeting specific criteria for handover preparation, the gNB-P sends configuration requests to each potential target cell (gNB-T_1 and gNB-T_2 in the Figure). The potential target cells reply with the configuration needed to either hop from gNB-P to gNB-T and from gNB-T to gNB-P (this helps to *i)* quickly recover from potential failures and *ii)* current gNB-P may be a future target cell). When the UE has received the configurations from potential target cells, upon meeting the criteria that define the conditionality of the handover execution, s/he can execute a HO to one among all the target cells that satisfy such criteria. 3GPP

does not mandate specific rules for such choices. Next, we will harness such a degree of freedom and narrow down the best target cell leveraging the fact that bus routes have an inherent predictable aspect.

5.2. FCHO for Public Transit System

We elaborated in Section 4.1 - Connectivity - that mobile devices attach preferentially to the same BSs over time because of pre-determined routes. Based on those findings, we expect that also handovers happen with preferential links between BSs. To verify such intuition, we perform the following analysis.

Hypergraph Structure. We formally denote $\mathcal{H} = (\mathcal{V}, \mathcal{E})$ as an unweighted directed hypergraph, where $\mathcal{H}$ is the set of vertices and $\mathcal{E}$ is the set of hyperedges. Let $n_v = |\mathcal{V}|$ be the number of vertices in $\mathcal{H}$ and $n_e = |\mathcal{E}|$ be the number of hyperedges in $\mathcal{H}$. We build one hypergraph per mobile network operator and per bus route over the course of the measurement period. Specifically, we assign each individual BS_i to a vertex $v_i \in \mathcal{V}$ and each handover between source BS_1 (i.e., v_1) and target BS_2 (i.e., v_2) becomes an directed edge $e_{1\to2}$. The hyperedge $E = \{e_{1\to2}, e_{2\to3}, \ldots, e_{n-1\to n}\}$ ultimately contains a sequence of BSs that the mobile device has attached to during the bus route. Each hyperedge $E \in \mathcal{E}$ represents one measurement day, hence, by construction, $\mathcal{E}$ denotes sequences of handovers of each bus route over time. We now construct an adjacency matrix $\mathcal{A}$ for each $\mathcal{H}$ with $\mathcal{A}$ being a square matrix in which rows and columns are indexed by the vertices of $\mathcal{H}$. For each $v_i, v_j \in \mathcal{V}$ with $v_i \neq v_j$, the elements of $\mathcal{A}$ are defined as $a_{v_i,v_j} = |e_{v_i\to v_j} \in \mathcal{E}|$.

Fig. 16 portrays for one urban and one suburban routes the adjacency matrices built on the corresponding hypergraphs for each operator. These matrices are derived over a measurement period of 9 or more days. We observe that in each matrix there exist elements a with occurrences way higher than all other elements in the respective row and columns: this indicates that handovers happen with preferential links. This behavior holds for both urban and suburban routes. Note that, by the construction of $\mathcal{A}$, there exists self-links a_{v_i,v_i} in the same vertex: this occurs because of cell re-selection. We highlight in red two examples in both

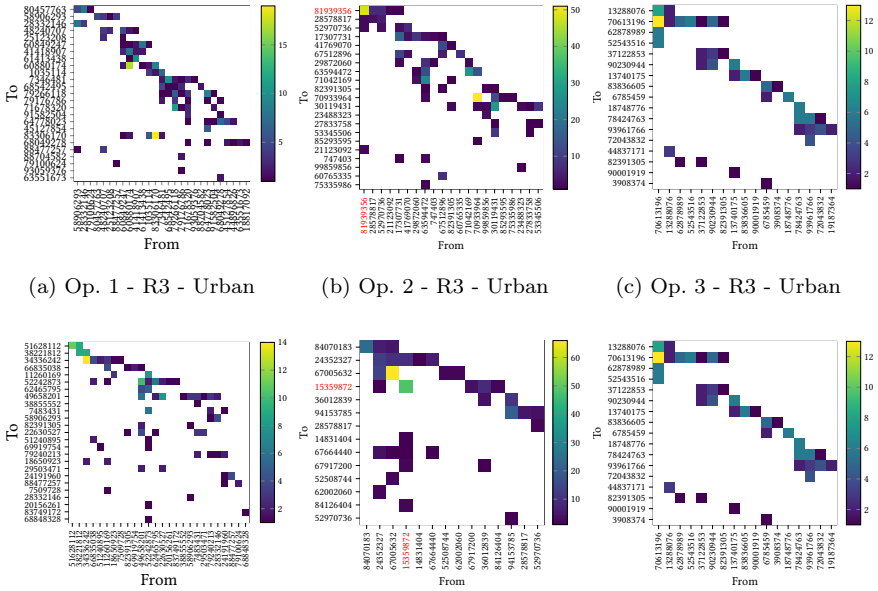

(a) Op. 1 - R3 - Urban (b) Op. 2 - R3 - Urban (c) Op. 3 - R3 - Urban

(d) Op. 1 - R2 - Suburban (e) Op. 1 - R2 - Suburban (f) Op. 3 - R3 - Urban

Figure 16: Scatter plot with handover occurrences between the source (from) and target (to) BSs in urban and suburban routes

urban and suburban routes respectively in Fig. 16(b) and in Fig. 16(e).

$\mathcal{H}$-based Target Cell Selection. Based on the above observations, let us now come back to the FCHO details. The network prepares up to 8 configurations for target gNBs, which is more than the cardinality of possible handovers in each row and column of the observed $\mathcal{A}$ in Fig. 16. As the 3GPP specifies that is the UE that decides on the target cell to hop to but does not mandate any rule for such choice, we argue that based on historical knowledge, i.e., $\mathcal{A}$, the UE can make a good choice by selecting for each source gNB the edge with the highest occurrence in the corresponding column in $\mathcal{A}$. Formally, we define the $\mathcal{H}$-based Target Cell Selection ($\mathcal{H}$-TCS) for each source $v_i \in \mathcal{V}$, the target gNB $v_j \in V$ is $e_{v_i \to v_j} = \text{argmax}_{v_i}\{\mathcal{A}_{v_i,v_j}\}$.

We now apply such a technique in a what-if-analysis that aims to identify the potential benefits of the FCHO in avoiding cell reselection (i.e., the elements $a_{v_i,v_i} \in \mathcal{A}$). Further, we ask ourselves which is a reasonable amount of time for building $\mathcal{H}$ so that the corresponding $\mathcal{A}$ is representative and can be utilized by

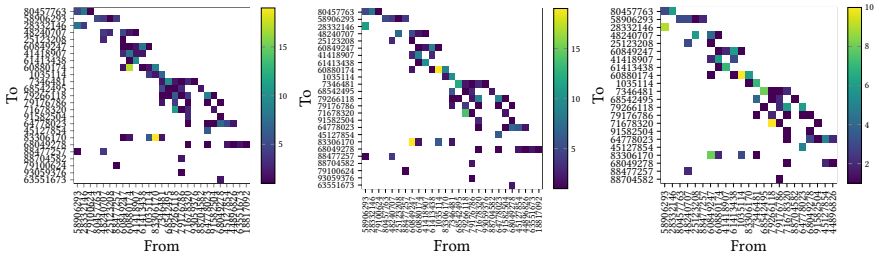

(a) $\mathcal{H}$-TCS not enabled (b) $\mathcal{H}$-TCS enabled (c) $\mathcal{H}$-TCS with 5 days observations

Figure 17: Comparison of handover occurrences in different FCHO scenarios for Operator 1

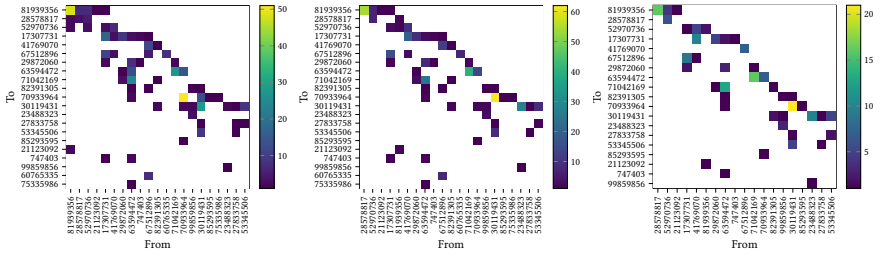

(a) $\mathcal{H}$-TCS not enabled (b) $\mathcal{H}$-TCS enabled (c) $\mathcal{H}$-TCS with 5 days observations

Figure 18: Comparison of handover occurrences in different FCHO scenarios for Operator 2

users for target cell selection. For this analysis, we set 5 days the observation period, which is reasonable because it represents the weekdays of one single week. We show the results for the urban route R3 in the figures Fig. 17, Fig. 18 and Fig. 19 for Op. 1, Op. 2, and Op. 3 respectively. In each figure, we compare the baseline case where we do not apply the hypergraph-based target cell selection (subfigures (a)), the case where we use the new target cell selection (subfigures (b)), and one case where we build $\mathcal{H}$ only with 5 days long observation periods (subfigures (c)). Across all operators, the figures show that the following observations hold: *(i)* enabling $\mathcal{H}$-TCS augments the magnitude of the highest occurrences that are not self-links, *(ii)* the 5 days long observation period is sufficient to characterize $\mathcal{A}$ and therefore helpful for target cell selection because all the highest occurrence links are captured only with lower magnitude.

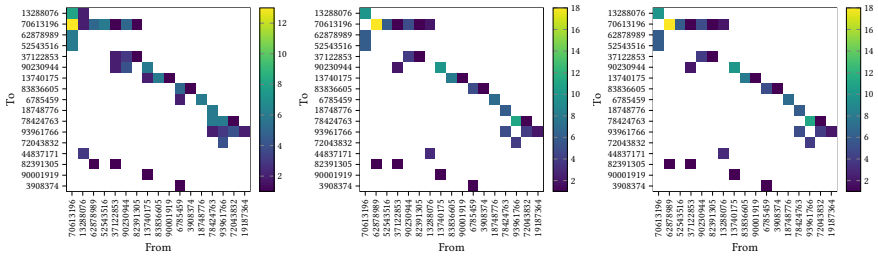

(a) $\mathcal{H}$-TCS not enabled (b) $\mathcal{H}$-TCS enabled (c) $\mathcal{H}$-TCS with 5 days observations

Figure 19: Comparison of handover occurrences in different FCHO scenarios for Operator 3

6. 5G Network Performance

In this section, we discuss the end-to-end network performance of commercial 5G networks when using mobile devices on the public bus transit system.

Methodology. Like previous studies [3, 4, 7] we use the Ookla's Speedtest application [31]. The application can operate in two modes: by default, it connects to a geographically close server, but it also allows to manually select the server to connect to. We leverage the latter mode for the latency tests and manually select servers in several Spanish cities where the network operators studied host servers (namely, Bilbao, Valencia, Sevilla, and Barcelona). For throughput and video streaming, we use the default app configuration. Note that for all the network indicators, Speedtest indicates the peak performance, i.e., the maximum attainable throughput or video resolution and the minimum attainable latency. While this does not correspond to the user-perceived network quality and experience, it allows our study to focus on understanding network performance in different transit scenarios (urban vs. suburban) with minimal impact from various factors like connection technology and network congestion. After having identified that there are no major performance differences between the M2001J2G and the SM-G988U phones, we selected the first one for our tests. The following results are obtained by repeating each test 10 times in several route parts (4 urban and 4 suburban - see § 2) on different days for two consecutive weeks. Unlike previous studies [7], our measurements are taken

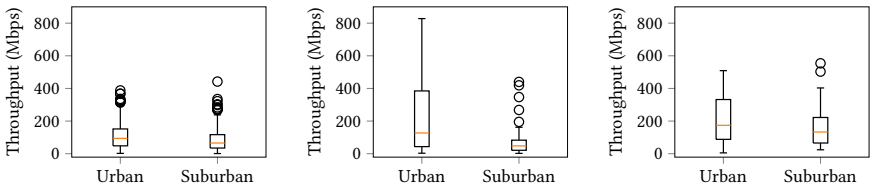

(a) Operator 1 (b) Operator 2 (c) Operator 3

Figure 20: Network throughput in urban and suburban scenarios

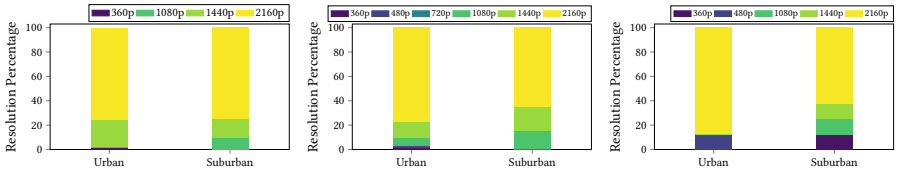

(a) Operator 1 (b) Operator 2 (c) Operator 3

Figure 21: Video resolution in urban and suburban areas

under mobility where was not possible to guarantee a clear line of sight with the gNB persistently across experiments.

Throughput. Fig. 20 shows the observed network throughput for all the network operators studied in urban and suburban routes. Note that we do not stress the network to fully saturate the available bandwidth, but we rather use the bandwidth as a normal user would do. We find that mobile connections in urban bus routes always achieve higher throughput than suburban ones (see the median of the boxplots). Op. 2's mobile network deployment in urban and suburban areas is comparatively the most different among all the studied operators and indicates poor 5G penetration outside the main city center area and excellent coverage in the city center (the interquartile range, i.e., the area between the 25th and 75th percentile of the urban class is the highest). The high throughput is attributed to the carrier bandwidth which is the highest configured among the three operators (see Table 2). On the contrary, Op. 1's deployment exhibits the lowest difference in attainable throughput in urban and suburban environments. Op. 3 positions itself in between Op. 2 and Op. 1 and the attainable throughput in suburban areas is the highest among all the operators. It should be noted that the median throughput of Op. 2 is higher than that of Op. 3.

Video Streaming Quality. Fig. 21 shows for the different operators the percentage of the highest video resolution achieved during the 15-second streaming. The highest resolution is 2160p and corresponds to 3840×2160 pixels, commonly known as 4k. By contrast, the lowest resolution is 360p, which corresponds to a 480×360 pixels. Note that Speedtest implements adaptive bitrate (ABR) technology, hence the instantaneous video resolution is likely to vary. With this test, our objective is not to understand which ABR algorithm works better, but rather which is the peak resolution that is achievable in each scenario and to compare the different operators. We find that Op. 1 achieves comparable peak resolution in both urban and suburban areas while for Op. 2 and Op. 3 the peak resolution occurs more often in urban areas.

Latency. Fig. 22 shows the network latency for all the operators considered in the study in several Spanish cities[4]. The server located in Madrid is always the one with the fastest response, which highlights the benefits of edge computing with servers located close to end users. We also identify zones that correspond to the domain of responsibility of each User Plane Function (UPF) that defines the border of the MNO network before the traffic is routed to the public Internet. For Op. 1 and Op. 2, Barcelona is likely to be in a different zone than the other cities (the median response is above 32 ms) while for Op. 3 we identify Bilbao.

7. Concluding Remarks

We have performed a comprehensive measurement study of 5G performance when moving on the public bus transit system of a major EU city, Madrid. With commercial state-of-the-art mobile phones, we define an app-based methodology to uncover the current status of the mobile network operator configurations. Further, we have studied coverage, mobility management, and end-to-end performance. Regarding mobility management, we have analyzed (fast) conditional handover solutions conceived for 5G NR and proposed a new target cell selection

[4]For Op. 3 we could not test the server in Barcelona.

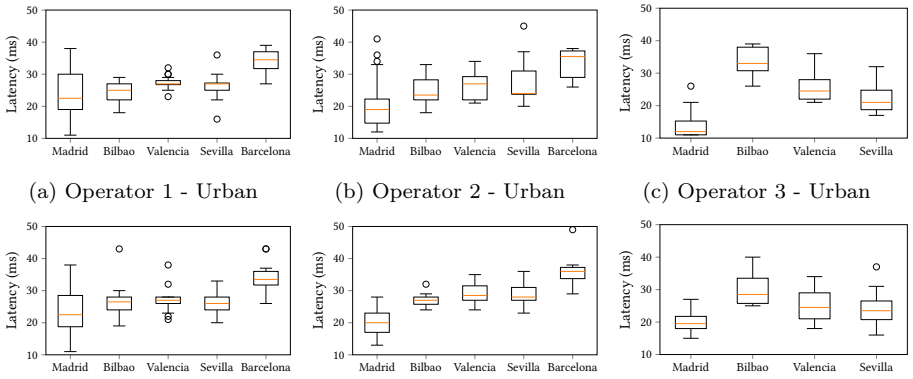

(a) Operator 1 - Urban (b) Operator 2 - Urban (c) Operator 3 - Urban

(d) Operator 1 - Suburban (e) Operator 2 - Suburban (f) Operator 3 - Suburban

Figure 22: Latency in urban and suburban areas

strategy that works well in the presence of predictable mobility like that of bus routes. Our three monthly long measurement campaign has revealed key aspects of the 5G operation in both city center and suburban municipalities. We have released our extensive dataset and the artifacts to the research community and plan to release the code for the analysis of the present submission upon acceptance.

Acknowledgment

This work is supported by the Juan de la Cierva Incorporation grant (IJC2019-039885-I) and grant PID2021-128250NB-I00 ("bRAIN") from the Spanish Ministry of Science and Innovation, and partially by the Madrid Regional Government through the TAPIR-CM program (S2018/TCS-4496), by the project RISC-6G, reference TSI-063000-2021-59, granted by the Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO-5G R&D Program of the Spanish Recovery, Transformation and Resilience Plan.

References

[1] C. Fiandrino, D. Juárez Martínez-Villanueva, J. Widmer, Uncovering 5G performance on public transit systems with an app-based measurement study,

in: Proc. of ACM MSWiM, 2022, p. 65–73. doi:10.1145/3551659.3559040.

[2] S. Neumeier, P. Wintersberger, A.-K. Frison, A. Becher, C. Facchi, A. Riener, Teleoperation: The holy grail to solve problems of automated driving? sure, but latency matters, in: Proc. of the ACM AUTOMOTIVEUI, 2019, p. 186–197. doi:10.1145/3342197.3344534.

[3] A. Narayanan, E. Ramadan, J. Carpenter, Q. Liu, Y. Liu, F. Qian, Z.-L. Zhang, A first look at commercial 5G performance on smartphones, in: Proc. of The Web Conference, 2020, p. 894–905. doi:10.1145/3366423.3380169.

[4] D. Xu, A. Zhou, X. Zhang, G. Wang, X. Liu, C. An, Y. Shi, L. Liu, H. Ma, Understanding operational 5G: A first measurement study on its coverage, performance and energy consumption, in: Proc. of ACM SIGCOMM, 2020, p. 479–494. doi:10.1145/3387514.3405882.

[5] D. Raca, D. Leahy, C. J. Sreenan, J. J. Quinlan, Beyond throughput, the next generation: A 5G dataset with channel and context metrics, in: Proc. of ACM MMSys, 2020, p. 303–308. doi:10.1145/3339825.3394938.

[6] A. Narayanan, E. Ramadan, R. Mehta, X. Hu, Q. Liu, R. A. K. Fezeu, U. K. Dayalan, S. Verma, P. Ji, T. Li, F. Qian, Z.-L. Zhang, Lumos5G: Mapping and predicting commercial MmWave 5G throughput, in: Proc. of ACM IMC, 2020, p. 176–193. doi:10.1145/3419394.3423629.

[7] A. Narayanan, X. Zhang, R. Zhu, A. Hassan, S. Jin, X. Zhu, X. Zhang, D. Rybkin, Z. Yang, Z. M. Mao, F. Qian, Z.-L. Zhang, A variegated look at 5G in the wild: Performance, power, and qoe implications, in: Proc. of ACM SIGCOMM, 2021, p. 610–625. doi:10.1145/3452296.3472923.

[8] E. Ramadan, A. Narayanan, U. K. Dayalan, R. A. K. Fezeu, F. Qian, Z.-L. Zhang, Case for 5G-aware video streaming applications, in: Proceedings of the 5G-MeMU, 2021, p. 27–34. doi:10.1145/3472771.3474036.

[9] M. I. Rochman, V. Sathya, N. Nunez, D. Fernandez, M. Ghosh, A. S. Ibrahim, W. Payne, A comparison study of cellular deployments in chicago and miami using apps on smartphones, in: Proc. of ACM WiNTECH, 2022, p. 61–68. doi:10.1145/3477086.3480843.

[10] A. Narayanan, M. I. Rochman, A. Hassan, B. S. Firmansyah, V. Sathya, M. Ghosh, F. Qian, Z. Zhang, A comparative measurement study of commercial 5G mmWave deployments, in: Proc. of IEEE INFOCOM, 2022, pp. 800–809.

[11] A. Hassan, A. Narayanan, A. Zhang, W. Ye, R. Zhu, S. Jin, J. Carpenter, Z. M. Mao, F. Qian, Z.-L. Zhang, Vivisecting mobility management in 5G cellular networks, in: Proc. of ACM SIGCOMM, 2022, p. 86–100. doi:10.1145/3544216.3544217.

[12] K. Kousias, M. Rajiullah, G. Caso, O. Alay, A. Brunstorm, L. De Nardis, M. Neri, U. Ali, M.-G. Di Benedetto, Coverage and performance analysis of 5G non-standalone deployments, in: Proc. of ACM WiNTECH, 2022, p. 61–68. doi:10.1145/3556564.3558233.

[13] R. A. K. Fezeu, E. Ramadan, W. Ye, B. Minneci, J. Xie, A. Narayanan, A. Hassan, F. Qian, Z.-L. Zhang, J. Chandrashekar, M. Lee, An in-depth measurement analysis of 5G mmwave PHY latency and its impact on end-to-end delay, in: Passive and Active Measurement, Springer Nature Switzerland, Cham, 2023, pp. 284–312.

[14] Y. Liu, C. Peng, A close look at 5G in the wild: Unrealized potentials and implications, in: Proc. of IEEE INFOCOM, 2023, pp. 1–10.

[15] P. Parastar, A. Lutu, O. Alay, G. Caso, D. Perino, Spotlight on 5G: Performance, device evolution and challenges from a mobile operator perspective, in: Proc. of IEEE INFOCOM, 2023, pp. 1–10.

[16] C. Lee, H. Cho, S. Song, J.-M. Chung, Prediction-based conditional handover

for 5g mm-wave networks: A deep-learning approach, IEEE Vehicular Technology Magazine 15 (1) (2020) 54–62. doi:10.1109/MVT.2019.2959065.

[17] S. B. Iqbal, S. Nadaf, A. Awada, U. Karabulut, P. Schulz, G. P. Fettweis, On the analysis and optimization of fast conditional handover with hand blockage for mobility (2023). doi:10.1109/ACCESS.2023.3260630.

[18] 3GPP TSG-RAN WG2 NR Adhoc Meeting #2, Conditional handover basic aspects and feasibility. r2-1706489., (R. 15) (2017).

[19] S. Kim, S. Lee, E. Ko, K. Jang, J. Yeo, Changes in car and bus usage amid the COVID-19 pandemic: Relationship with land use and land price, Journal of Transport Geography 96 (2021) 103168. doi:https://doi.org/10.1016/j.jtrangeo.2021.103168.

[20] J. English, Why public transportation works better outside the U.S., Bloomberg CityLab - Transportation - `https://tinyurl.com/2m9fcyhs`, accessed on 2022-05-09 (2018).

[21] S. Hemminki, P. Nurmi, S. Tarkoma, Accelerometer-based transportation mode detection on smartphones, in: Proc. of the ACM SenSys, 2013, p. 1–14. doi:10.1145/2517351.2517367.

[22] Y. Liao, J. Gil, R. H. Pereira, S. Yeh, V. Verendel, Disparities in travel times between car and transit: Spatiotemporal patterns in cities, Scientific reports 10 (1) (2020) 1–12.

[23] S. B. Damsgaard, N. J. Hernández Marcano, M. Nørremark, R. H. Jacobsen, I. Rodriguez, P. Mogensen, Wireless communications for internet of farming: An early 5g measurement study, IEEE Access 10 (2022) 105263–105277. doi:10.1109/ACCESS.2022.3211096.

[24] S. Rameshselvakumar, T. Arjun, T. Mohan Kumar, C. Naveen Kumar, Teleoperated campus utility vehicle using SDAS based on CAN, in: Proc. of IEEE ICAECA, 2021, pp. 1–6. doi:10.1109/ICAECA52838.2021.9675639.

[25] H. Elsherbiny, A. M. Nagib, H. Abouzeid, H. M. Abbas, H. S. Hassanein, A. Noureldin, A. Bin Sediq, G. Boudreau, 4G LTE network data collection and analysis along public transportation routes, in: Proc. of the IEEE GLOBECOM, 2020, pp. 1–6. doi:10.1109/GLOBECOM42002.2020.9348031.

[26] Y. Pan, R. Li, C. Xu, The first 5G-LTE comparative study in extreme mobility, Proc. ACM Meas. Anal. Comput. Syst. 6 (1) (Feb 2022). doi:10.1145/3508040.

[27] Y. Li, C. Peng, Z. Zhang, Z. Tan, H. Deng, J. Zhao, Q. Li, Y. Guo, K. Ling, B. Ding, H. Li, S. Lu, Experience: A five-year retrospective of mobileinsight, in: Proc. of the ACM MobiCom, 2021, p. 28–41. doi:10.1145/3447993.3448138.

[28] A. Narayanan, E. Ramadan, J. Quant, P. Ji, F. Qian, Z.-L. Zhang, 5g tracker–a crowdsourced platform to enable research using commercial 5g services, in: Poster at ACM SIGCOMM, 2020, pp. 65–67.

[29] GNetTrack, Application, accessed 03/06/2022 (2022).
URL https://gyokovsolutions.com/g-nettrack/

[30] Application, Network signal guru., accessed 03/06/2022 (2022).
URL https://www.qtrun.com/en/

[31] OOKLA, Speedtest for android., accessed 03/06/2022 (2022).
URL https://www.speedtest.net/apps/android

[32] A. Mahimkar, et al., Auric: Using data-driven recommendation to automatically generate cellular configuration, in: Proc. of ACM SIGCOMM, 2021, p. 807–820. doi:10.1145/3452296.3472906.

[33] N. Ludant, G. Noubir, SigUnder: A stealthy 5g low power attack and defenses, in: Proc. of ACM WiSec, 2021, p. 250–260. doi:10.1145/3448300.3467817.

[34] 3GPP TS38.213, Physical layer procedures for control., (R. 16.9) (2022).

[35] L. Meegahapola, T. Kandappu, K. Jayarajah, L. Akoglu, S. Xiang, A. Misra, Buscope: Fusing individual & aggregated mobility behavior for "live" smart city services, in: Proc. of ACM MobiSys, 2019, p. 41–53. doi:10.1145/3307334.3326091.

[36] H. Martikainen, I. Viering, A. Lobinger, T. Jokela, On the basics of conditional handover for 5G mobility, in: Proc. IEEE PIMRC, 2018, pp. 1–7. doi:10.1109/PIMRC.2018.8580946.

[37] Ericsson, Mobility Report, June 2023. Technical Report., Accessed on 04/07/2023: `https://www.ericsson.com/en/reports-and-papers/mobility-report/reports/june-2023` (2023).

[38] 3GPP TS 23.401, General packet radio service (GPRS) enhancements for evolved universal terrestrial radio access network (e-utran) access., (R. 16.12) (2022).

[39] 3GPP TS 23.502, Procedures for the 5G system (5GS)., (R. 16.13) (2022).

[40] Samsung, 4g-5g interworking, accessed on 20/06/2022: `https://tinyurl.com/yvn4wnuy` (2017).

[41] 3GPP TS 38.331, Radio resource control (RRC) protocol specification., (R. 16.18) (2022).

[42] J. Stanczak, U. Karabulut, A. Awada, Conditional handover in 5G - principles, future use cases and FR2 performance, in: Proc. of IWCMC, 2022, pp. 660–665. doi:10.1109/IWCMC55113.2022.9824571.

[43] E. Juan, M. Lauridsen, J. Wigard, P. Mogensen, Performance evaluation of the 5G NR conditional handover in LEO-based non-terrestrial networks, in: Proc. of IEEE WCNC, 2022, pp. 2488–2493. doi:10.1109/WCNC51071.2022.9771987.

[44] Z. Huang, X. Xu, T. Lu, H. Li, C. Sun, S. Wang, Mobile integrated access and backhaul node backhaul link failure recovery with

conditional handover, in: Proc. of IEEE ICISCAE, 2022, pp. 418–422. doi:10.1109/ICISCAE55891.2022.9927593.

[45] 3GPP TSG RAN Meeting #94e, New wid on further nr mobility enhancements, document rp-213565, Potential target (Rel-18) (2021).

[46] Z.-H. Huang, Y.-S. Chen, M.-J. Tsai, Efficient conditional handover algorithm in 5G with blockages using recurrent neural network, in: Proc. of IEEE CCNC, 2023, pp. 686–687. doi:10.1109/CCNC51644.2023.10059816.

[47] A. Prado, H. Vijayaraghavan, W. Kellerer, ECHO: Enhanced conditional handover boosted by trajectory prediction, in: Proc. of IEEE GLOBECOM, 2021, pp. 01–06. doi:10.1109/GLOBECOM46510.2021.9685348.

Biographies

Claudio Fiandrino is a Senior Researcher at IMDEA Networks Institute. He obtained his Ph.D. degree from the University of Luxembourg in 2016. His current research interests encompass AI applicability to 5G and 6G mobile networks, including network optimization, explainability and trustworthiness. Claudio was awarded with a Fulbright/José Castillejo mobility grant, two Spanish Juan de la Cierva grants, and several Best Paper Awards for his research.

David Juárez Martínez-Villanueva is a M.Sc. student at University Carlos III of Madrid, Spain. He spent at IMDEA Networks 6 months performing his B.Sc. final work during 2022.

Joerg Widmer is Research Professor and Research Director of IMDEA Networks in Madrid, Spain. His research focuses on wireless networks, in particular millimeter-wave communications. Joerg Widmer is an IEEE Fellow and Distinguished Member of the ACM, and was awarded an ERC consolidator grant, the Friedrich Wilhelm Bessel Research Award of the Alexander von Humboldt Foundation, a Mercator Fellowship of the German Research Foundation, and a Spanish Ramon y Cajal grant.