Multi-Operator Connectivity Sharing for Reliable Networks: A Data-Driven Risk Analysis

André Gomes, Jacek Kibiłda, *Senior Member, IEEE*, Arman Farhang, Ronan Farrell, *Member, IEEE*, and Luiz A. DaSilva, *Fellow, IEEE*

Abstract—A key distinction between today's and future networks is the appetite for reliable communication to support emerging critical-communication services. In this paper, we study multi-operator connectivity as a form of redundancy to support the design of reliable networks and investigate its trade-offs. This approach is motivated by 3GPP standardisation initiatives of dual-connectivity and similar techniques in industrial wired networks. We deploy a risk awareness performance metric to assess reliability: this superquantile metric accounts for periods of connectivity shortfalls. Our analysis shows that multi-operator connectivity brings significant reliability gains, in particular when network deployments by different operators exhibit high complementarity in coverage. We also explore the effects of multi-connectivity on spectral efficiency in times of high demand for bandwidth. Our study is based on a real-world dataset comprising signal strength indicators of three mobile operators in Dublin, Ireland.

Index Terms—Network reliability, network sharing, multi-operator connectivity, multi-connectivity, risk analysis.

I. INTRODUCTION

RELIABLE communication is one of the grand challenges for the next generation of mobile networks, enabling emerging communication services for factory automation, augmented reality, cloud gaming, smart transportation, and other yet-to-come applications [1], [2]. Reliability usually requires redundant network resources, such as antennas and spectrum, or denser network deployments, for example, in the form of additional base stations (BSs). Both approaches result in additional investment by mobile network operators (MNOs).

Network sharing is an alternative for mobile operators to mitigate upfront investments while leveraging more network resources being made available when needed. This may come in different flavors such as neutral hosting, where operators share a common pool of network resources and infrastructure [3], or multilateral agreements, where operators make network resources available to each other's subscribers [4].

We have previously studied network sharing in the form of multi-operator connectivity sharing [5], where a mobile can simultaneously connect to multiple network operators. This approach is motivated by the increasing number of mobiles equipped with multi- and embedded-SIM (eSIM) cards, as well as initiatives to standardize multi-connectivity since 3GPP release 12. Furthermore, combinations of deployments by different operators exhibit spatial characteristics (such as clustering [6]) that may provide additional support for reliable connectivity through increased redundancy.

We focus on multi-operator connectivity as a form of redundancy, where information is redundantly transmitted in all active connections, in light of 3GPP standardisation initiatives of dual-connectivity since release 12 and similar approaches in use in industrial wired networks for reliable communication [7]. Our prior work [5], alongside others [8], [9], indicates that multi-operator connectivity can be an efficient way to achieve reliable communication. The gains mostly lie in periods of connectivity shortfalls, when multi-connectivity is better able to meet the aggregate demand and avoid service outages. However, redundant connections come at a cost to operators in terms of lower spectral efficiency. This is especially harmful in high-demand networks where the use of redundant resources introduced by sharing can lead to increasing demand for spectrum (or any other scarce network resource), potentially decreasing the capacity operators can offer to their subscribers.

In this paper, we extend our previous work [5] and conduct an investigation of the gains, costs, and trade-offs of multi-operator connectivity for reliable communication. To that end, we first revisit how network performance is conventionally assessed and make the case for an alternative means of measuring reliability that captures periods of connectivity shortfalls. Then, we make use of a real-world dataset of signal quality measurements for three mobile operators in the city of Dublin, Ireland, to quantify the impacts of multi-operator connectivity, adopting the proposed reliability metric.

The remainder of the paper is organized as follows. Sec. II discusses related works and our contributions. Sec. III introduces the system model, the proposed metric for network reliability, and the data we use to study the impacts of multi-operator connectivity on reliable communication. Sec. IV

presents our data analysis, followed by final remarks and conclusions in Sec. V.

II. RELATED WORK AND CONTRIBUTIONS

In this section, we provide an overview of related works and our main contributions. Multi-connectivity has received attention as a solution to enhance network performance. A prominent example is the use of multiple connections to increase aggregate average performance. For instance, [10], [11], [12], [13] are variants of traditional network protocols to support multi-connectivity; however, in such cases, optimizing for mean performance often comes at the expense of reliability (when viewed as the suppression of outlier performance) as discussed in [13].

In other studies, multi-connectivity is an alternative approach to realism reliable communication – for comprehensive surveys, see [14], [15]. ReMP TCP [16] extends the multipath transmission control protocol (MPTCP) [10] to support reliable communication through the duplication of packets over multiple connections. The proposed approach suppresses the tail of the latency distribution to a greater extent than MPTCP and transmission control protocol (TCP) in both system-level simulations and experimental evaluation of dual-connectivity in an LTE mobile operator. In [8], the authors leverage diverse radio technologies (namely WiFi, 2G, 3G, and 4G) to enhance network reliability. They explore several approaches to encode and split information over multiple connections. In practical scenarios, packet duplication significantly reduces latency while enhancing reliable communication, outperforming other more sophisticated packet splitting strategies. In [9], the authors implement multi-base station connectivity for reliable communication in a small-scale testbed. Their approach requires coherent signal aggregation at the physical layer (PHY), which makes it challenging to implement in large-scale real-world scenarios because of stringent synchronization requirements between BSs. In [17], the authors study multi-user networks in multi-connectivity scenarios. They propose a matching algorithm to assign secondary connections so as to optimism data-rate subject to minimum performance requirements. The authors of [18], [19] deploy multi-connectivity in two use cases, respectively: (a) an experimental testbed at the Hamburg seaport where multi-connectivity enhances network reliability by suppressing disruptions during handovers for mobile barges equipped with sensing devices; and (b) a small-scale industrial hall where multi-connectivity decreases the outage probability for mobile devices. [20] focuses on the use of multi-connectivity in wireless LAN, whereas [21] discusses architectural enhancements of LTE and 5G New Radio for multi-connected mobiles, both in light of reliable communication.

Multi-connectivity often requires that more network resources are made available (in the form of antennas, spectrum, network density, etc.) to support multiple connections, which necessitates additional investments by mobile operators. Network sharing is an alternative way that mobile operators can leverage additional network resources while restricting capital and operational expenditures. The gains often come from the complementarity of multiple operators: there is no point in sharing resources if load and deployment patterns are the same [22], [23]. The analysis in [22] indicates that combining two operators can increase effective capacity in a cost-effective manner. The authors of [23] compare different forms of network sharing: (a) capacity sharing, modelled as a roaming process between two MNOs; (b) spectrum sharing, where bandwidth surplus is redirected to heavily loaded BSs; and (c) virtuality sharing, where spectrum is locally shared within mobile virtual networks. Their findings report on the effectiveness of capacity sharing, which performs better and is simpler to implement than the other forms of sharing that they studied. In [24], the potential of infrastructure and spectrum sharing is studied in the light of specific spatial characteristics of multi-operator networks, such as clustering, which is often found in real-world deployments [6]. On the operational side of network sharing, [25] shows that MNOs can achieve significant energy savings by jointly switching off BSs in multi-operator networks while guaranteeing quality-of-service (QoS) to users. Allowing operators to jointly manage their networks can also increase effectiveness by curbing idle capacity while minimizing associated expenditures [26]. The collaboration versus competition dilemma is an important question in network sharing economics, and [27] studies whether two MNOs should share network infrastructure. The authors demonstrate that network sharing has the potential to yield gains in some situations even if a MNO has the power to suppress the other or under competition regulation.

On the one hand, the aforementioned papers address multi-connectivity for increased network reliability and network sharing for cost-effective provisioning of network resources. On the other hand, none explores the potential of both approaches being jointly deployed. We took a step further in [5] and studied multi-operator connectivity sharing as a method to realism reliable communication. Our prior findings are focused on the coverage gains derived from multi-operator connectivity and their impact on network reliability. Here, we revisit these results in Sec. IV-A and extend our analysis by considering multi-user finite-bandwidth networks under high demand. As we shall see subsequently, assigning multiple connections to multiple users can overload the network and eventually compromise the network performance. Our new study reports on the associated trade-offs based on real-world data analysis and provides a set of lessons on how operators can leverage multi-operator connectivity sharing for reliable communication while minimizing the loss of aggregate capacity.

A. Main Contributions

Our major contributions are fourfold. First, we extend our previous study in [5] by considering multi-user scenarios with shared network resources. Second, we consider the case of network sharing in the form of multi-operator connectivity as a means to enhance network reliability. Multiple connections are considered as redundant resources and packets are duplicated over the active connections at higher network layers rather than PHY-based approaches. Third, we adopt superquantiles as a

measure of reliability, for (as we shall discuss in Sec. III-C) it captures the severity of periods of impaired connectivity, unlike traditional performance metrics. Fourth, we base our analysis on a real-world dataset, which is, to the best of our knowledge, unique because it comprises simultaneous measurements of signal strength indicators of three mobile networks taken during a walk-test campaign.

III. PRELIMINARIES

A. Network Redundancy

Reliability often requires that redundant resources are made available. This work focuses on multi-connectivity as a form of redundancy: multiple simultaneously active connections are used to transmit redundant information. This can be achieved in different ways. The parallel redundancy protocol (PRP) introduced the concept of packet duplication for industrial Ethernet networks [7], where every node is connected to two distinct and independent networks that transmit duplicated packets. Similarly, [7] also defines the high-availability seamless redundancy protocol (HSR), which uses a ring network topology to provide redundancy. Both approaches are attractive for reliable communication as they are simple and robust fault tolerant solutions and require zero time recovery in the occurrence of failures in either connection. More recently, PRP-like packet duplication was also introduced in mobile networks since 3GPP release 12, with the concept of dual-connectivity (DC). Redundancy can be handled at the packet data convergence protocol (PDCP) layer, and each connection functions with independent medium access control (MAC) and PHY layers, requiring no signal aggregation nor tight time synchronization between network entities. Carrier aggregation (CA) is another alternative for providing redundant network resources by making use of multiple component carriers to transmit information. Unlike DC, however, CA uses cells within the same cell group and a single MAC entity [28]. In coordinated multi-point (CoMP), different transmission entities (different BSs or even multiple antennas on the same BS) jointly coordinate communication to a mobile [29] and can be used to improve reliability.

We address reliable communication through the adoption of multi-operator connectivity to provide network redundancy. CA, CoMP, or general MAC/PHY-based solutions require significant coordination and tight time synchronization between BSs. This is especially challenging to implement across multiple independently operated networks. Our focus in this paper is on PRP-like packet duplication in light of its simplicity and robustness as well as 3GPP standardisation initiatives in DC. Packets are duplicated at the PDCP layer as shown in Fig. 1. Connections are independent MAC and PHY entities, requiring no signal aggregation at the PHY layer nor coordination between BSs.

B. Network Model

We consider a network with M operators and N mobile users, where each mobile is embedded with multiple radio frequency (RF) front-ends, each of which can support simultaneous connection to a different MNO. All the mobiles are


Fig. 1. Example of multi-operator connectivity sharing with packet aggregation at the PDCP layer.

connected to their native operator, and some are also allowed to multi-connect to some other operators in compliance with what their native operator permits. The multi-operator connections are held as in Fig. 1, which exemplifies a shared network with two operators where the mobile is granted a secondary connection to MNO_1 so that it can still receive information reliably, even in case of signal blockage or link outage from its native connection to MNO_0.

Reliability is ultimately considered end-to-end; however, we focus on the radio access network (RAN) because the wireless environment is often the most dynamic, limiting, and expensive component of mobile networks. We consider the downlink communication and assess the system's performance based on the effective channel capacity in use by the mobile. We assume each operator occupies its own licensed spectrum and thus operators do not interfere with each other. Each transmitted packet is duplicated in all active connections and aggregated at the PDCP layer. Redundant copies of packets by a mobile are discarded. We assume network interfaces are independent hardware-based implementations with negligible overhead impact on network performance. The effective channel capacity of a mobile n connected to a set of base stations, $\mathcal{B}_n$, is given by:

$$z_n = \max\{w_{n,b} \times \log_2(\gamma_{n,b} + 1), \forall b \in \mathcal{B}_n\} \text{ bps}. \quad (1)$$

where $w_{n,b}$ and $\gamma_{n,b}$ are the bandwidth and signal-to-noise ratio (SNR) of the link between mobile n and base station b.

C. Network Reliability

Networks are often assessed by their average performance. However, the introduction of mission-critical communication services poses stringent reliability demands to be met by the network. These services are averse to the risk of performance fluctuations, such as periods of connectivity shortfalls that are unlikely to be captured by a simple average. In this section, we look at alternative metrics that depict the risk of under-performance as measures of reliability.

Let Z_m be a random variable of a figure of merit (e.g., channel capacity or spectrum efficiency) of a typical subscriber of mobile network operator m. The conventional way of evaluating reliability is the use of quantizes as a risk measure. The α-quantile, where $\alpha \in [0, 1]$ corresponds to the associated degree of risk, captures information about the tail of the

Fig. 2. The CDF of two dummy random variables as an example of variables with equal 0.1-quantiles but distinct distributions.

density function of Z_m. The formal definition is:

$$q_\alpha(Z_m) = \sup\{z \in \mathbb{R} : F_{Z_m}(z) \leq \alpha\}. \qquad (2)$$

where $F_{Z_m}(z)$ is the cumulative distribution function of Z_m.

Quantiles provide a means for quantitative comparison of distinct networks and network designs in terms of their ability to meet minimum design goals. Nevertheless, the use of quantiles brings some concerns. Figures of merit for different networks may have the same quantiles at a degree α, i.e., $q_\alpha(Z_o) = q_\alpha(Z_p)$, even though they may have distinct density functions (see Fig. 2 for instance). This is especially critical if the underlying differences are in the lower-tail of the corresponding densities, for quantiles are insensitive to the remaining low-probability events beyond their values. In practical deployments, this is important because services may somewhat tolerate under-performance by running fallback routines and prioritizing the transmission of highly critical information, depending on how severe such events of impaired connectivity are. Therefore, quantiles fail to address questions such as *what is the risk associated with exceeding the desirable requirements?* and *what to expect in the worst case?* Another concern is that quantiles lack desirable mathematical properties, such as convexity, which limits their ease of use in design optimization (refer to [30] for further discussion).

Similarly to quantiles, superquantiles are a measure of the tail of a distribution, but, unlike quantiles, they account for the magnitude of the excess. They were proposed by [30], [31] as a measure of risk; superquantiles tend to be more mathematically tractable than quantiles, being especially useful in heavy-tailed distributions where quantiles do not consider the magnitude of lower-probability events. They correspond to the expectation over the lower-tail[1] of the density function of a random variable, and also depend on the parameter α that corresponds to the associated degree of risk. The formal definition is as follows:

$$\bar{q}_\alpha(Z_m) = \frac{1}{\alpha} \int_0^\alpha q_\beta(Z_m) d\beta. \qquad (3)$$

One interesting mathematical property of superquantiles is that they approach the expected value when $\alpha \to 1$ and the infimum when $\alpha \to 0$, depicting both the conventional risk-neutral and the ultimate risk-averse approaches in the extreme cases.

[1]Or upper-tail, depending on the figure of merit of interest.

To illustrate the usefulness of such a metric, let us consider a critical remote control service that requires a performance indicator of $R_{\min}$ to be fully operational but somewhat supports under-performance by running fallback routines. A real-life example is [32], where the remote-controlled robotic arm can extrapolate next movements from previous commands but halts depending on the severity of impaired connectivity. Let Z_0 and Z_1 be the performance indicators of two mobile networks over time, where greater values indicate greater performance. We can set $\alpha = 1$ and estimate the expected performance from the networks over time and compare it against $R_{\min}$. In Fig. 2, the average performance of Z_1 (approx. 0) favors its choice in comparison with Z_0 (approx. -2). If the system is required to be fully operational 90% of the time, we can then pose further questions such as *what is the risk associated with the 10% worst-case scenarios?* If risk is measured by quantiles, both Z_0 and Z_1 perform similarly in our example, and network 1 would be a preferable choice based on its greater average performance; however, notice that Z_1 has a heavier lower-tail than Z_0. By using superquantiles as a measure of risk, the remote controller can estimate the risk and *severity* associated with such cases (Z_0: $\bar{q}_{\alpha=0.1}$ approx. -3.6; Z_1: $\bar{q}_{\alpha=0.1}$ approx. -37) and (if possible) devise plan B routines accordingly. This approach empowers the service controller with a means of ranking different networks based on the risk and severity each of them offers at a degree α. The same reasoning applies to contrast network designs and rank their respective degrees of dependability.

In this paper, we conduct a risk assessment on the reliability of connectivity to mobile networks. For the aforementioned reasons, we make use of superquantiles to analyst the gains of multi-operator connectivity sharing for the design of reliable networks.

D. Our Data

We conducted walk tests and collected measurements from three mobile operators in Dublin, Ireland (we will refer to them as MNO_0, MNO_1, and MNO_2). The resulting data includes timestamps, geographical coordinates, and performance metrics such as signal-to-noise ratio (SNR), reference signal received power (RSRP), and reference signal received quality (RSRQ) as shown in Tab. I. These measurements were collected using G-MoN, a freeware passive observation application [33]. For a fair comparison, we used three LTE mobile phones (one per operator) of the same model and brand. The use of G-MoN involved activating the application to record data to a script file. During the walk tests, the phones were encased in a frame in a backpack to keep them at a similar orientation to each other so as to limit bias across the different operators. The phones recorded the data from the BS to which they were attached, once every second. In open areas of heavy traffic, the walkers would walk to a small number of separate points and pause for two to five minutes to gather a larger number of data samples. There was no downtime, i.e., the phones were always connected to a base station. Fig. 3 depicts our traces over three different geographical areas of the city, namely North Dublin, the region around the Guinness Storehouse, and South Dublin.
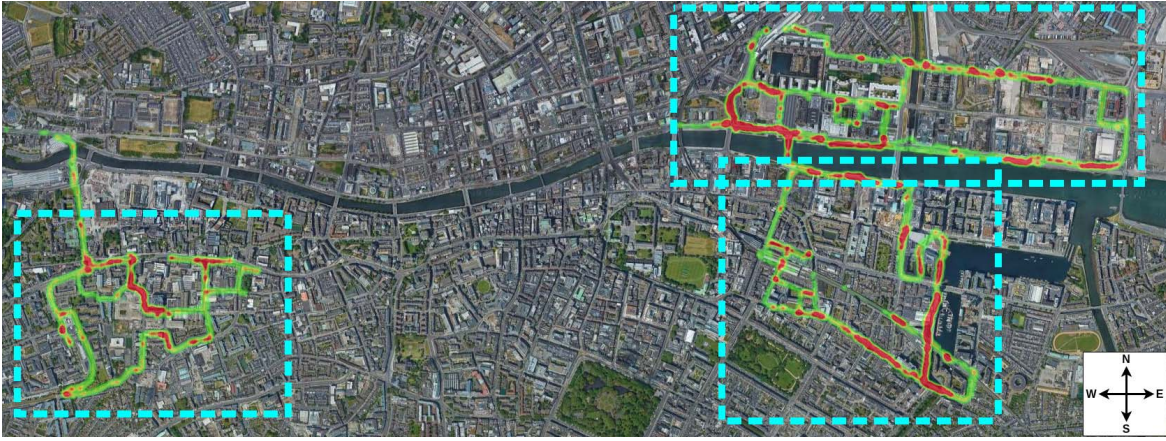
Fig. 3.　Signal-to-noise ratio (SNR) traces of an operator for locations (from left-to-right, top-to-bottom): (a) North Dublin, (b) Guinness storehouse, and (c) South Dublin. Color legend: ● < -10 dB, -10 dB $\leq$ ● < 0 dB, and 0 dB $\leq$ ●.

TABLE I
ILLUSTRATIVE EXAMPLE OF OUR TRACES WITH THREE KEY
PERFORMANCE INDICATORS, NAMELY SNR, RSRP, AND RSRQ, PER
OPERATOR AT SAME TIMES AND LOCATIONS

Time	Coord	MNO_0	MNO_1	MNO_2
t_0	$coord_0$	γ_0,rsrp$_0$,rsrq$_0$	γ_0,...	γ_0,rsrp$_0$,rsrq$_0$
t_1	$coord_1$	γ_1,rsrp$_1$,rsrq$_1$	γ_1,...	γ_1,rsrp$_1$,rsrq$_1$
...	...	...	...	...
t_i	$coord_i$	γ_i,rsrp$_i$,rsrq$_i$	γ_i,...	γ_i,rsrp$_i$,rsrq$_i$

TABLE II
ILLUSTRATIVE EXAMPLE OF THE FINAL DATASET

| User id	Coord	MNO_0		MNO_1		MNO_2	
		BS id	SNR	BS id	SNR	BS id	SNR
u_0	$coord_0$	b_0	$\gamma_{0,0}$	b_1	$\gamma_{0,1}$	b_2	$\gamma_{0,2}$
...	...	...	...	...	...	...	...
u_j	$coord_j$	b_a	$\gamma_{j,a}$	b_b	$\gamma_{j,b}$	b_c	$\gamma_{j,c}$

Fig. 4.　*Coverage gains.* The lowest 10-percentile of SNR of the Guinness storehouse micro-region of our data.

IV. DATA ANALYSIS

In this section, we look at how multi-operator connectivity impacts mobile communication. We assess the downlink channel capacity as a figure of merit for the three operators in our dataset. We contrast multi-operator connectivity (MC) with single-operator connectivity (SC) so that the gains and losses of MC are made explicit. Subscripts indicate the MNOs in use. For example, we denote SC_0 to indicate that subscribers of MNO_0 operate in SC, MC_{01} to indicate that subscribers of MNO_0 and MNO_1 multi-connect to both operators' BSs (dual-connectivity), and MC_{012} to indicate that subscribers of the three operators multi-connect to BSs of each other (triple-connectivity). The total available bandwidth at each base station is set to 1 MHz for simplicity. We ran 100 Monte Carlo experiments for each network scenario (i.e., single-operator connectivity and variants of multi-operator connectivity), each experiment consisting of N_m active subscribers for each mobile operator m, where N_m is a Poisson random variable of mean λ_m. Let λ be the average density of simultaneously active mobiles per BS and k_m the number of BSs of MNO_m, such that $\lambda_m = \lambda k_m$. The N_m active users are randomly selected from our dataset and assigned to the operator's network for each Monte Carlo experiment. The placement of active mobiles is prone to overcrowding in a few BSs, for

Unique as they are, our traces have two main shortcomings to represent the network model of Sec. III-B. First, they are limited to a single user per operator. We circumvent this limitation by considering that each entry in Tab. I could have been produced by a distinct user. The areas of heavy traffic have a higher density of data points. This mimics the geographical density of users and reflects locations where there is more demand for network resources. Second, there is no information on which BS each mobile was connected to at the time of data collection. In the absence of this information and operators' specific association policies, we assume that each user is connected to the geographically closest BS. To that end, we deploy a second public dataset on top of our traces [34], which contains the geographical coordinates of BSs of the three operators of interest. While in reality there may be cases when the mobile does not associate with the closest BS, sporadic mismatches do not significantly impact our analysis; our model and data capture the geographical demand of users for connectivity and the placement of network equipment to serve them. The resulting modified dataset is structured as shown in Tab. II and consists of 24763 possible users/locations and 40 base stations of MNO_0, 52 of MNO_1, and 56 of MNO_2.

(a) Channel capacity of MNO_0 subscribers.

(b) Channel capacity of MNO_1 subscribers.

(c) Channel capacity of MNO_2 subscribers.

(d) The α-superquantile of CC for MNO_0.

(e) The α-superquantile of CC for MNO_1.

(f) The α-superquantile of CC for MNO_2.

(g) The PDF of CC for MNO_0 subscribers.

(h) The PDF of CC for MNO_1 subscribers.

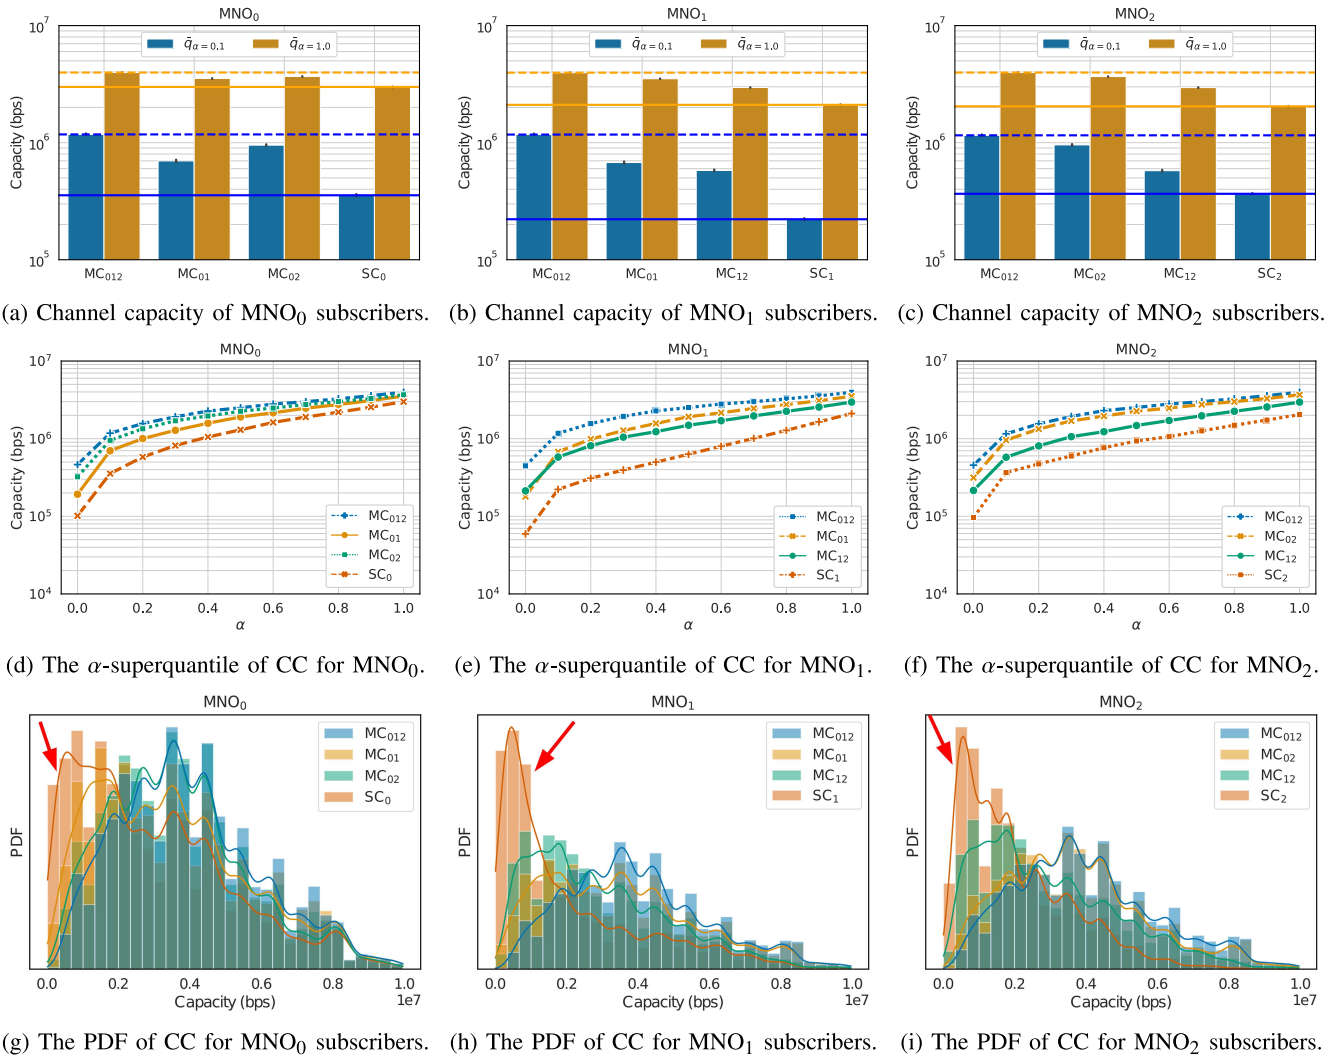(i) The PDF of CC for MNO_2 subscribers.

Fig. 5. *Coverage gains*. Channel capacity (CC). All mobiles multi-connect in MC.

some locations comprise higher density of data points as mentioned in Sec. III-D. In real-world deployments, users would be eventually blocked once resources are finished until operators deploy more BSs in such areas to counteract overcrowding. In our experiments, we assume sufficient network infrastructure is already deployed to meet geographical demand, and therefore, subject each BS to a maximum of 100 active subscribers. This corresponds to the maximum number of resource blocks in an LTE system with 20 MHz bandwidth divided into resource blocks 180 kHz-wide and 10 kHz guard-bands.

A. Coverage Gains

Intuitively, the gains from multi-operator connectivity come from the complementary coverage and demand of different operators. That is, one operator may provide extra network resources when another fails. This is specially true for signal coverage because of the distinct placement of base stations, consequently leading to different signal propagation patterns. We reinforce this intuition by plotting the 10%-lowest SNR data points of a micro-region of our data in Fig. 4. Each color refers to an operator: MNO_0 (red), MNO_1 (blue), and MNO_2

(yellow). In the area highlighted by the dashed-circle, all three operators have weak signal coverage. Interestingly, however, the remainder of the data points indicates a high degree of complementary coverage between operators, such as pointed by the white arrow, where red predominates.

We start our analysis by evaluating the coverage gains of multi-operator connectivity. In MC_{ab}, all subscribers of MNO_a multi-connect to MNO_b's network and vice-versa. Let us consider the channel capacity, a density λ of 10 active mobiles per BS, and the same unit bandwidth is allocated to all mobiles so that $w_{n,b} = 1$ MHz $\forall n \in N$ in Eq. (1).

Figs. 5(a)–5(c) show the 0.1-superquantile (blue bars) as well as the average capacity (orange bars) per operator. The interesting information is the relative gain from SC to MC. The horizontal lines highlight this difference, where continuous lines represent the performance of SC and dashed lines refer to the best performance of MC. Operators benefit the most from multi-connectivity when all three operators participate in the shared network, i.e., MC_{012}. By contrasting the difference between continuous and dashed lines, we can clearly notice that the gains are much higher for the 0.1-superquantile

than for the average user. Recall that the $\bar{q}_{\alpha=0.1}$ refers to the lower-tail of the distribution of the capacity. The difference in gains decreases as we move from the lower-tail ($\alpha \to 0$) to the upper-tail ($\alpha \to 1$) as shown in Figs. 5(d)–5(f). This supports the application of multi-operator connectivity for reliable communication, as the gains mostly reside in situations where operators alone under-perform. These results corroborate the observation from Fig. 4 that operators tend to complement each other in scenarios of weak signal coverage.

The same results can be viewed from another angle as shown in Figs. 5(g)–5(i), the empirical probability density function (PDF)[2] of the capacity for the subscribers of each of the three operators, under SC and MC. The most remarkable difference between single- and multi-operator connectivity lies in the lower-tail of the distributions (see red arrows). The smaller areas under the lower tails of the PDF for the case of MC explain the greater gains in the 0.1-superquantile that we have previously discussed. This is intuitive, as operators often plan their network to provide overall good signal coverage in the same areas to supply their similar customers' demands, such as in high-density business locations (refer to [6] for clustering in multi-operator networks). The difference usually resides in parcels that are hard to plan for because, for example, of signal propagation effects or limitations given the location of the base stations.

To confirm this intuition, we explore the correlation of the coverage among operators. That is, if an MNO has strong/weak signal coverage, how likely is it that neighboring operators also offer good/weak coverage? We pose this question with respect to all data points in our dataset and also regarding the intersection of the 10%-lowest SNR data points of each mobile operator. Figs. 6(a) and 6(b) show the Pearson correlation coefficients concerning the two cases. One can notice a significant smaller correlation for the 10-percentile of the SNR data points, indicating that operators exhibit stronger complementarity in locations of weak signal strength than overall. A remark is that dual-connectivity combinations that lead to gains close to MC_{012} in Figs. 5(a)–5(c) also present the smallest correlation coefficients, MC_{02} (dual-connectivity between MNO_0 and MNO_2), MC_{01}, and MC_{02}, respectively. This indicates, as we expect from intuition, that combining MNOs that strongly complement each other effectively yields coverage gains.

The results we have discussed rely on the assumption that each mobile in the network has the same fraction of the bandwidth $w_{n,b}$, making the capacity a function of the SNR only. This might relate to network deployments where spectrum is abundant. In reality, however, spectrum is often scarce and expensive, and the use of redundant connections in MC tends to increase its demand.

B. Spectrum in Demand

In this section, we analyst a similar scenario to Sec. IV-A but consider a frequency-division multiple access (FDMA)

(a) All mobiles.   (b) The 10-percentile.

Fig. 6. *Coverage gains*. Pearson correlation coefficients of SNR regarding all data points and the 10-percentile.

network whereby the bandwidth available at each BS is allocated in a round-robin fashion and equally serves all the active mobiles, such that $w_{n,b} = 1/(\sum_{\forall n' \in b} 1)$ MHz for all active mobiles connected to each base station b. Hence, the effective channel capacity in Eq. (1) depends on the number of active mobiles connected to each BS as well as on the SNR experienced by a mobile.

Figs. 7(a)–7(c) are the counterparts of Figs. 5(a)–5(c) and show the gains from SC to MC in terms of the 0.1- and 1.0-superquantiles. Again, the horizontal lines contrast the best performance of MC with SC for each MNO. We see some reliability gains, especially for MNO_1, but the gains are modest in comparison with the previous case. This is even worse for the average capacity, which decays from SC to MC indicating a loss of performance.

In Figs. 7(d) and 7(e), we further investigate such results by analyzing the cost of multi-operator connectivity. As redundant connections increase spectrum use, we look into the spectrum efficiency (SE), the ratio between the effective channel capacity and the total bandwidth allocated for each mobile n. Both SC and MC lead to similar spectrum efficiency in the worst case (i.e., $\bar{q}_{\alpha=0.1}$), whereas MC is up to half as efficient as SC on average. The spectrum efficiency observed by mobile user n is calculated as:

$$s_n = \frac{z_n}{\sum_{\forall b \in \mathcal{B}_n} w_{n,b}} \; [\; bps/Hz]. \tag{4}$$

Recall that the channel capacity is a linear function of the bandwidth and logarithmic in the SNR, so competition for spectrum can lead to lower capacity for individual mobiles, despite coverage gains. However, we have previously pointed out that the coverage gains are modest on average, implying that MC may be ineffective in scenarios of high demand for network resources.

We take a step further in our analysis and quantify the benefit that multi-operator connectivity brings to each user. To that end, we track the connection with the highest achievable capacity. If that connection is from the mobile's native operator, secondary links are not being used and are a waste of allocated bandwidth because the home operator already provides the best radio channel for communication. Otherwise, a mobile benefits from MC, as secondary links outperform the primary connection.

In Fig. 8(a), each data point represents the SNR that a mobile experiences from its native operator. For ease of

[2]The histograms correspond to the density of capacity measures from our experiments, whereas continuous lines are the kernel density estimate of the PDF.
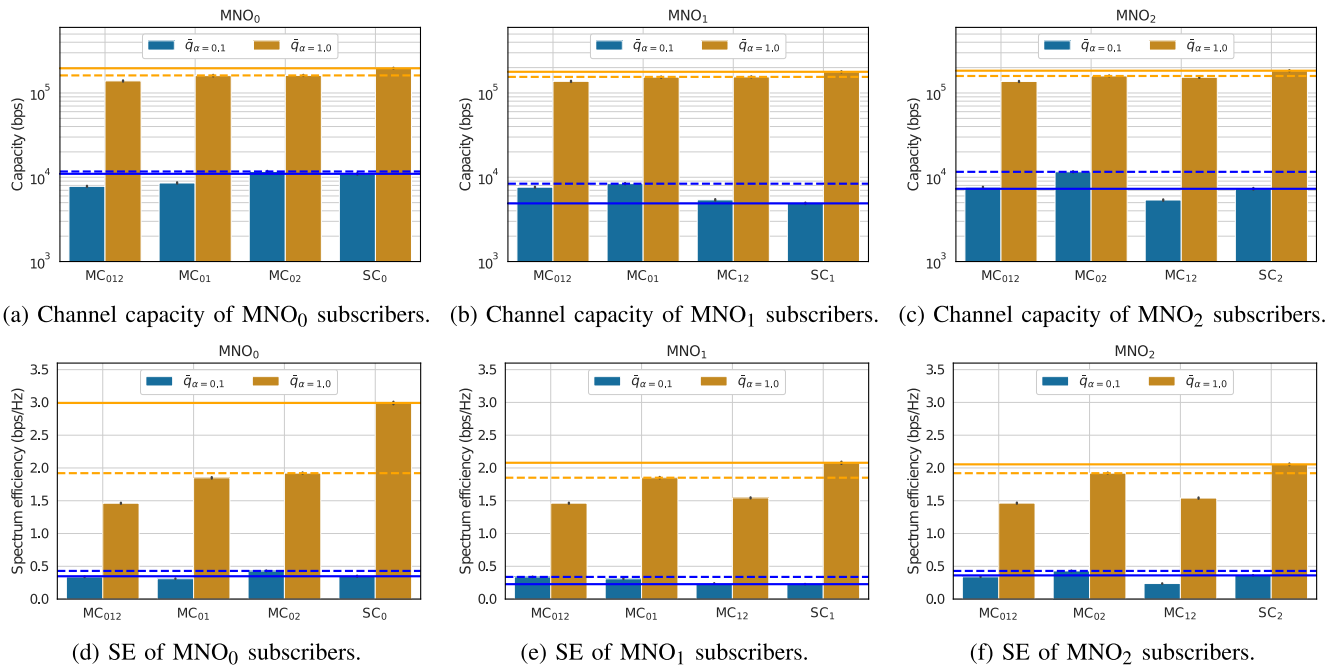
(a) Channel capacity of MNO_0 subscribers.   (b) Channel capacity of MNO_1 subscribers.   (c) Channel capacity of MNO_2 subscribers.

(d) SE of MNO_0 subscribers.   (e) SE of MNO_1 subscribers.   (f) SE of MNO_2 subscribers.

Fig. 7. *Spectrum in demand*. Channel capacity (CC) and spectrum efficiency (SE). All mobiles multi-connect in MC.

(a) Strip plot of beneficiaries.
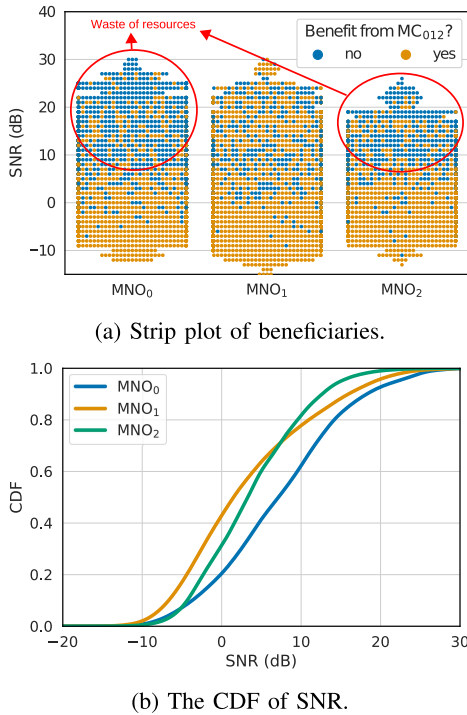

(b) The CDF of SNR.

Fig. 8. *Spectrum in demand*. Beneficiaries of MC and empirical cumulative distribution function (CDF) according to SNR.

visualization, we randomly limit the number of data points presented. The orange dots represent users that have benefited from MC (we refer to the benefit from full multi-operator connectivity, i.e., MC_{012}), whereas the blue dots those that have not. There are two remarkable facts that we can qualitatively observe here. First, many mobiles do not benefit from MC – approx. 46% for MNO_0, 23% for MNO_1, and 33% for MNO_2. Second, the mobiles that benefit the most

are the ones under weak signal coverage from their native operators – notice the higher density of orange dots for low values of SNR. The higher density of beneficiaries with high values of SNR for MNO_1 is because of its weaker signal coverage in comparison with others, lagging behind other operators in approximately 70% of our data as shown in Fig. 8(b). As a consequence, MNO_1 is more prone to link outages and a prominent candidate to benefit from multi-operator connectivity.

Motivated by our observation that mobiles experiencing low SNR are likely to benefit from MC, we analyst a scenario where we only allow the subscribers of the (lower) β-percentile of SNR of each operator to multi-connect, where $\beta \in [0, 1]$. Those users are designated by their native operator and based on the instantaneous SNR of their native connection. The remaining subscribers are only granted a connection to their native operators. Let us first focus on the subscribers of MNO_0. Figs. 9(a) and 9(d) show the 0.1- and 1.0-superquantiles of the channel capacity, respectively. multi-operator connectivity (MC) has a significant increase in reliability for small values of β (reaching up to 50.9% gains in MC_{02}) while having a low impact on the average performance. As more mobiles multi-connect ($\beta \to 1$), the gains proportionally diminish, impacting both 0.1- and 1.0-superquantiles of the channel capacity.

These results stem from the impact of MC on spectrum efficiency. For small values of β, the spectrum efficiency of MC is even higher than SC as shown in Figs. 9(g) and 9(j). As we have previously pointed out, the users of low SNR are likely to benefit from MC. However, the spectrum efficiency decreases as mobiles that are already well covered by their native operator are allowed to also maintain redundant connections to other operators, leading to performance losses.

(a) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_0 subscribers. (b) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_1 subscribers. (c) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_2 subscribers.

(d) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_0 subscribers. (e) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_1 subscribers. (f) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_2 subscribers.

(g) The $\bar{q}_{\alpha=0.1}$ of SE of MNO_0 subscribers. (h) The $\bar{q}_{\alpha=0.1}$ of SE of MNO_1 subscribers. (i) The $\bar{q}_{\alpha=0.1}$ of SE of MNO_2 subscribers.

(j) The $\bar{q}_{\alpha=1.0}$ of SE of MNO_0 subscribers. (k) The $\bar{q}_{\alpha=1.0}$ of SE of MNO_1 subscribers. (l) The $\bar{q}_{\alpha=1.0}$ of SE of MNO_2 subscribers.
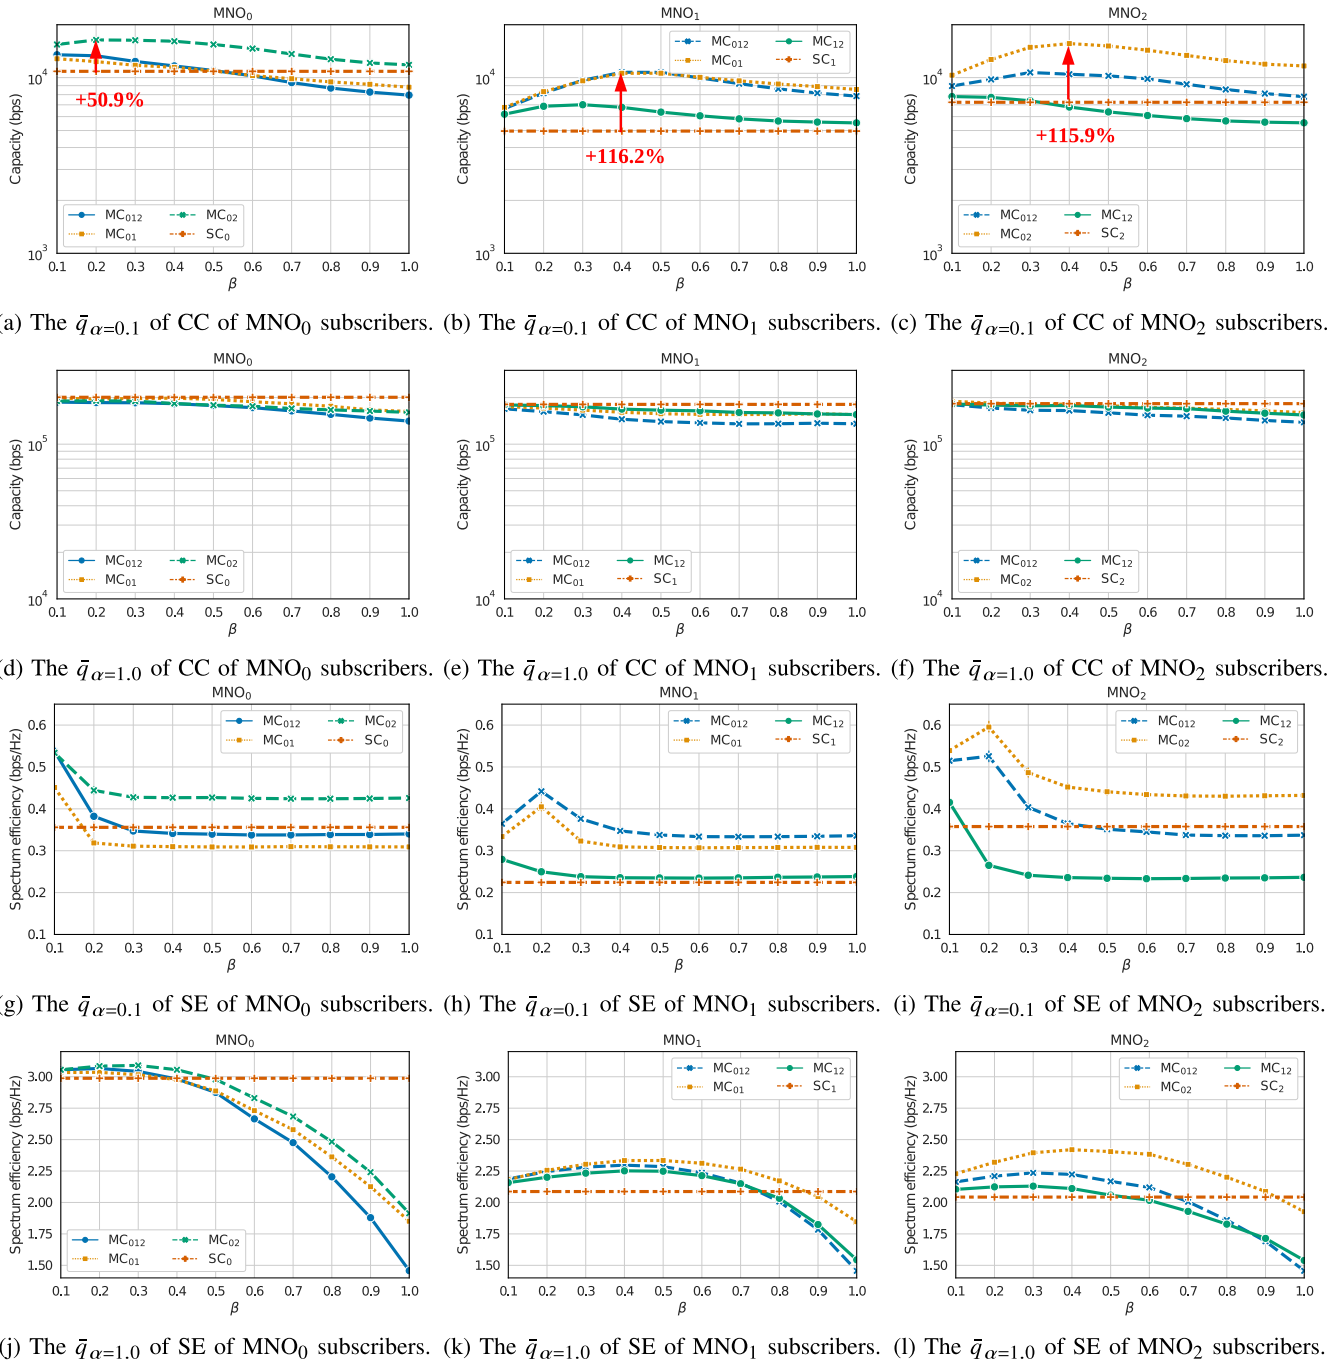
Fig. 9. *Spectrum in demand*. 0.1- and 1.0-superquantiles of the channel capacity (CC) and spectrum efficiency (SE). Only the subscribers of the lower β-percentile of SNR of each MNO multi-connect in MC; the remaining users single-connect as in SC.

The same trends are also observed for MNO_1 and MNO_2, as shown in Figs. 9(b) and 9(c) and Figs. 9(h) and 9(i). The only exception relates to subscribers of MNO_2 operating in MC_{12} in Fig. 9(i), which significantly underperforms SC_2 for a wide range of values of β. Recall from Fig. 6(b) that MNO_1 and MNO_2 are the only ones to exhibit positive correlation regarding their 10-percentile of SNR data points. The weak complementary in coverage indicates that operators are unlikely to benefit from each other in such a case, explaining the lower spectrum efficiency.

Interestingly, dual-connectivity is at least as good as MC_{012}: MNO_0's subscribers benefit the most from MC_{02} in Fig. 9(a); MNO_1's subscribers similarly benefit the most from MC_{012} and MC_{01} in Fig. 9(b); and MNO_2's subscribers benefit the most from MC_{02} in Fig. 9(c). A similar trend is shown in Figs. 9(d)–9(f), where dual-connectivity incurs smaller penalties than MC_{012} in average. This is good news because, as we have previously discussed, multiple connections often come at the expense of lower spectrum efficiency. The best dual combinations (MC_{02}, MC_{01}, and MC_{02} in Figs. 9(a)–9(c), respectively) coincide with small correlation coefficients in

(a) Channel capacity of MNO_0 subscribers. (b) Channel capacity of MNO_1 subscribers. (c) Channel capacity of MNO_2 subscribers.

Fig. 10.   *Scheduling*. The α-superquantile of the channel capacity for round-robin (RR) and proportional fair (PF) scheduling algorithms. $\beta = 0.3$, and $\lambda = 10$.

(a) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_0 subscribers. (b) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_1 subscribers. (c) The $\bar{q}_{\alpha=0.1}$ of CC of MNO_2 subscribers.

(d) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_0 subscribers. (e) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_1 subscribers. (f) The $\bar{q}_{\alpha=1.0}$ of CC of MNO_2 subscribers.

Fig. 11.   *Mobile density*. The channel capacity (CC) as a function of the average network density λ (active mobiles/BS). $\beta = 0.3$, and scheduling is RR.

C. Scheduling and Mobile Density

So far we have analyzed settings where the bandwidth is allocated in a round-robin (RR) fashion. In this section, we also consider bandwidth assignment according to the proportional fair (PF) scheduling strategy, which assigns bandwidth to each mobile user n according to the metric $w_{n,b} = c_{n,b}/(\sum_{\forall n' \in b} c_{n',b})$ MHz where $c_{n,b} = 1/\log_2(\gamma_{n,b}+1)$ is the weight associated with the SNR. In this way, more bandwidth is assigned to active mobile users that experience low SNR.

Fig. 10 shows the α-superquantile of the channel capacity of each mobile operator for RR and PF as a function of α. We set $\beta = 0.3$, as it yielded significant gains in reliability with a small penalty in terms of spectral efficiency, as discussed

in Sec. IV-B, and $\lambda = 10$. Let us first focus on Fig. 10(a), which depicts the channel capacity of subscribers of MNO_0. Therein, PF SC_0 outperforms RR SC_0 for small values of α. This is expected since PF favors mobile users of poor SNR, allocating more bandwidth to them. For $\alpha > 0.2$, in turn, RR SC_0 tends to outperform PF SC_0.

Similar patterns are also observed for the different combinations of MC. Interestingly, PF MC does not yield significantly higher gains than RR MC for small values of α. At first glance, we would expect a significant increase in reliability in PF MC as both proportional fair and multi-operator connectivity tend to benefit under-performing users. However, it is noticeable that PF and MC are independently deployed in different network layers (MAC and PDCP, respectively). We study this effect in more detail by considering conflicts between PF and MC. A conflict is when the interface of the highest allocated bandwidth (scheduled by PF) does not coincide with the interface of the highest achievable capacity (selected by MC). We denote conflict-ratio as the percentage

of mobiles that experience conflict between PF and MC. Here, the conflict-ratios of MNO_0, MNO_1, and MNO_2 are approximately 50%, 70%, and 65%, respectively, for multi-connected mobiles in PF MC_{012}. This result is intriguing as it indicates that conflicting network protocols can compromise (rather than jointly enhance) network reliability.

Lastly, we look at one more facet of multi-operator connectivity sharing, its dependence on the number of mobiles. As we have discussed in the preceding sections, the increasing demand for network resources can counterbalance the coverage gains of MC. Recall that N_m, for all operators $m \in M$, is modelled as a truncated Poisson random variable dependent on the average density λ of simultaneously active mobiles per BS. Intuitively, this is an important aspect of MC as the number of active mobiles in the network directly maps to the demand for spectrum. As we are about see, however, it affects both SC and MC to a similar extent.

Figs. 11(a) and 11(d) show the results for the 0.1- and 1.0-superquantiles of the capacity as a function of the density of active mobiles of MNO_0. We set $\beta = 0.3$ and consider the round-robin scheduling strategy. As the density of active mobiles grows, so does the impact on the performance of SC and MC. Nonetheless, the difference between SC and MC is mostly consistent over different mobile densities, and the gains and losses of multi-operator connectivity loosely depend on the number of active mobiles in the network.

Although we have constrained our attention to the subscribers of MNO_0, the same trends are observed for other mobile operators, as shown in Figs. 10 and 11. The same applies for other values of $\beta \in [0, 1]$, and we have omitted those plots from this manuscript.

V. CONCLUSION

The next generation of mobile networks will face different challenges from today's. A paramount distinction is the increasing demand for reliable connectivity by critical-communication services. In this paper, we studied the gains, costs, and trade-offs of network sharing, in the form of multi-operator connectivity, as a means of achieving reliable communication. To that end, we made use of a real-world dataset of three mobile operators in the city of Dublin.

We explored how multi-operator connectivity impacts reliability. On the one hand, the coverage gains mostly reside in locations of weak signal strength, where operators tend to complement each other. On the other hand, the demand for limited network resources introduced by multiple users may throttle coverage gains and affect performance.

Mobile operators can increase reliability by focusing on the mobiles of weak SNR. Our findings indicate that a simple but effective strategy of (only) allowing weak SNR mobiles to multi-connect leads to a significant increase in reliability at a small loss of average capacity. It is also shown that the benefits of multi-operator connectivity are directly related to the complementary coverage of different operators.

The density of active users similarly affects single- and multi-operator connectivity. However, at a fixed density, the effectiveness of multi-operator connectivity decreases as more

mobiles multi-connect. Scenarios of high-reliability demand, where (almost) no user can afford under-performance, require other strategies for effectively allocating network resources in lieu of universal multi-connectivity. Additionally, our results suggest that conflict between scheduling strategies and MC can limit the gains of multi-operator connectivity. These motivate further investigation on how, when, and where to activate secondary interfaces (or to schedule network resources from multi-operator networks) to effectively realism reliable communication, which we shall pursue in our future works.

REFERENCES

[1] S. Chandrashekar, A. Maeder, C. Sartori, T. Höhne, B. Vejlgaard, and D. Chandramouli, "5G multi-RAT multi-connectivity architecture," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC)*, 2016, pp. 180–186.

[2] P. Popovski *et al.*, "Wireless access for ultra-reliable low-latency communication: Principles and building blocks," *IEEE Netw.*, vol. 32, no. 2, pp. 16–23, Mar./Apr. 2018.

[3] T. Li and L. Bai, "Model of wireless telecommunications network infrastructure sharing & benefit-cost analysis," in *Proc. Int. Conf. Inf. Manag. Innovat. Manag. Ind. Eng.*, vol. 2, 2011, pp. 102–105.

[4] T. Fiatal, "Mobile virtual network operator," U.S. Patent 8 107 921, Jan. 2012.

[5] A. Gomes, J. Kibiłda, A. Farhang, R. Farrell, and L. A. DaSilva, "Network sharing for reliable networks: A data-driven study," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2020, pp. 1–6.

[6] J. Kibiłda, B. Galkin, and L. A. DaSilva, "Modelling multi-operator base station deployment patterns in cellular networks," *IEEE Trans. Mobile Comput.*, vol. 15, no. 12, pp. 3087–3099, Dec. 2016.

[7] *Industrial Communication Networks—High Availability Automation Networks—Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)*, IEC Standard 62439-3:2012, 2012.

[8] J. J. Nielsen, R. Liu, and P. Popovski, "Ultra-reliable low latency communication using interface diversity," *IEEE Trans. Commun.*, vol. 66, no. 3, pp. 1322–1334, Mar. 2018.

[9] A. Wolf, P. Schulz, M. Dörpinghaus, J. C. S. S. Filho, and G. Fettweis, "How reliable and capable is multi-connectivity?" *IEEE Trans. Commun.*, vol. 67, no. 2, pp. 1506–1520, Feb. 2019.

[10] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP extensions for multipath operation with multiple addresses," Internet Eng. Task Force, Fremont, CA, USA, RFC 6824, Jan. 2013. [Online]. Available: http://www.rfc-editor.org/rfc/rfc6824.txt

[11] J. R. Iyengar, P. D. Amer, and R. Stewart, "Concurrent multipath transfer using SCTP multihoming over independent end-to-end paths," *IEEE/ACM Trans. Netw.*, vol. 14, no. 5, pp. 951–964, Oct. 2006.

[12] T. Viernickel, A. Froemmgen, A. Rizk, B. Koldehofe, and R. Steinmetz, "Multipath QUIC: A deployable multipath transport protocol," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2018, pp. 1–7.

[13] B. Walker, V. A. Vu, and M. Fidler, "Multi-headed MPTCP schedulers to control latency in long-fat/short-skinny heterogeneous networks," in *Proc. 13th Workshop Challenged Netw.*, 2018, pp. 47–54. [Online]. Available: https://doi.org/10.1145/3264844.3264847

[14] M. Suer, C. Thein, H. Tchouankem, and L. Wolf, "Multi-connectivity as an enabler for reliable low latency communications—An overview," *IEEE Commun. Surveys Tuts.*, vol. 22, no. 1, pp. 156–169, 1st Quart., 2020.

[15] G. J. Sutton *et al.*, "Enabling technologies for ultra-reliable and low latency communications: From PHY and MAC layer perspectives," *IEEE Commun. Surveys Tuts.*, vol. 21, no. 3, pp. 2488–2524, 3rd Qaurt., 2019.

[16] A. Frommgen, T. Erbshäußer, A. Buchmann, T. Zimmermann, and K. Wehrle, "ReMP TCP: Low latency multipath TCP," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2016, pp. 1–7.

[17] M. Simsek, T. Hößler, E. Jorswieck, H. Klessig, and G. Fettweis, "Multiconnectivity in multicellular, multiuser systems: A matching-based approach," *Proc. IEEE*, vol. 107, no. 2, pp. 394–413, Feb. 2019.

[18] P. Rost *et al.*, "Customized industrial networks: Network slicing trial at Hamburg seaport," *IEEE Wireless Commun.*, vol. 25, no. 5, pp. 48–55, Oct. 2018.

[19] A. H. Mahdi, T. Hößler, N. Franchi, and G. Fettweis, "Multi-connectivity for reliable wireless industrial communications: Gains and limitations," in *Proc. IEEE Wireless Commun. Netw. Conf. (WCNC)*, 2020, pp. 1–7.

[20] N. Schwarzenberg, A. Wolf, N. Franchi, and G. Fettweis, "Quantifying the gain of multi-connectivity in wireless LAN," in *Proc. Eur. Conf. Netw. Commun. (EuCNC)*, 2018, pp. 16–20.

[21] J. Rao and S. Vrzic, "Packet duplication for URLLC in 5G: Architectural enhancements and performance analysis," *IEEE Netw.*, vol. 32, no. 2, pp. 32–40, Mar./Apr. 2018.

[22] P. Di Francesco, F. Malandrino, and L. A. DaSilva, "Mobile network sharing between operators: A demand trace-driven study," in *Proc. ACM SIGCOMM Workshop Capacity Sharing*, 2014, pp. 39–44.

[23] J. S. Panchal, R. D. Yates, and M. M. Buddhikot, "Mobile network resource sharing options: Performance comparisons," *IEEE Trans. Wireless Commun.*, vol. 12, no. 9, pp. 4470–4482, Sep. 2013.

[24] J. Kibiłda, N. J. Kaminski, and L. A. DaSilva, "Radio access network and spectrum sharing in mobile networks: A stochastic geometry perspective," *IEEE Trans. Wireless Commun.*, vol. 16, no. 4, pp. 2562–2575, Apr. 2017.

[25] M. Oikonomakou, A. Antonopoulos, L. Alonso, and C. Verikoukis, "Cooperative base station switching off in multi-operator shared heterogeneous network," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, 2015, pp. 1–6.

[26] P. D. Francesco, F. Malandrino, T. K. Forde, and L. A. DaSilva, "A sharing- and competition-aware framework for cellular network evolution planning," *IEEE Trans. Cogn. Commun. Netw.*, vol. 1, no. 2, pp. 230–243, Jun. 2015.

[27] M. Andrews, M. Bradonjić, and I. Saniee, "Quantifying the benefits of infrastructure sharing," in *Proc. 12th Workshop Econ. Netw. Syst. Comput.*, 2017, pp. 1–6. [Online]. Available: https://doi.org/10.1145/3106723.3106734

[28] "Evolved universal terrestrial radio access (E-UTRA) and evolved universal terrestrial radio access network (E-UTRAN); Overall description, version 16.3.0," 3GPP, Sophia Antipolis, France, Rep. 36.300, 2020.

[29] D. Lee *et al.*, "Coordinated multipoint transmission and reception in LTE-advanced: Deployment scenarios and operational challenges," *IEEE Commun. Mag.*, vol. 50, no. 2, pp. 148–155, Feb. 2012.

[30] R. T. Rockafellar and J. O. Royset, "Engineering decisions under risk averseness," *ASCE-ASME J. Risk Uncertainty Eng. Syst. A, Civil Eng.*, vol. 1, no. 2, Jun. 2015, Art. no. 4015003.

[31] R. T. Rockafellar and S. Uryasev, "Conditional value-at-risk for general loss distributions," *J. Bank. Finan.*, vol. 26, no. 7, pp. 1443–1471, 2002.

[32] *Franka Emika: Minimum System and Network Requirements.* Accessed: Jan. 16, 2021. [Online]. Available: https://frankaemika.github.io/docs/requirements.html

[33] C. Knuetter. *G-MoN.* Accessed: Jul. 27, 2020. [Online]. Available: https://play.google.com/store/apps/details?id=de.carknue.gmon2

[34] *The World's Largest Open Database of Cell Towers.* Accessed: Jan. 2020. [Online]. Available: https://opencellid.org/

Arman Farhang received the Ph.D. degree from Trinity College Dublin, Ireland, in 2016. He was a Research Fellow with CONNECT, Trinity College Dublin, Ireland, from 2016 to 2018. He was a Lecturer with the School of Electrical and Electronic Engineering, University College Dublin, Ireland, for a short period in 2018. He joined the Department of Electronic Engineering, Maynooth University as a Lecturer in September 2018. He also holds an Adjunct Professor Position with TCD. He is an Associate Investigator with CONNECT Centre, where he leads the research around the topic of waveforms for 5G and beyond. He has published over 50 peer-reviewed publications on the topic of waveforms. His research interests include wireless communications, digital signal processing for communications, multiuser communications, and multiantenna and multicarrier systems. He serves as an Associate Editor of *EURASIP Journal on Wireless Communications and Networking*. He is a member of the organization committee of the IEEE Communication Society Flagship Conference ICC 2020. He has also served as a TPC member of many well-reputed IEEE conferences and workshops.

Ronan Farrell (Member, IEEE) received the Ph.D. degree in electronic engineering from University College Dublin, Dublin, Ireland, in 1998. Since 2000, he has been with Maynooth University, Maynooth, Ireland, where he is currently a Professor with the Department of Electronic Engineering. He is a Co-Principal Investigator with the Science Foundation Ireland CONNECT Centre for the Internet of Things. He leads a research team in the area of high-frequency radio systems with a particular interest in infrastructure systems, power amplifiers, and applications of millimeter wave communications. He has coauthored over 200 journal and conference papers.

André Gomes received the B.Sc. degree in telecommunications engineering from the Universidade Federal de São João del-Rei, Brazil, and the M.Sc. degree in computer science from the Universidade Federal de Minas Gerais, Brazil. He is currently pursuing the Ph.D. degree in electronic and electrical engineering with the Irish National Telecommunications Research Centre (CONNECT), Trinity College Dublin, Ireland. He was a visiting student with the University of Adelaide, Australia. His research interests include network reliability, self-organising networks, and software-defined networks.

Jacek Kibiłda (Senior Member, IEEE) received the Ph.D. degree from Trinity College, The University of Dublin, Ireland, in 2016. He is a Research Fellow with Trinity College Dublin and a Challenge Research Fellow with Science Foundation Ireland. He is also a Fulbright Scholar with The University of Texas at Austin from 2017 to 2018. His research focuses, broadly speaking, on understanding performance and design trade-offs in future mobile and wireless networks, using tools of stochastic geometry, optimization, and computer modeling.

Luiz A. DaSilva (Fellow, IEEE) is the Executive Director of the Commonwealth Cyber Initiative and a Professor with the Bradley Department of Electrical and Computer Engineering, Virginia Tech. He previously held the chair of Telecommunications with Trinity College Dublin, where he served as the Director of CONNECT, the Telecommunications Research Centre funded by the Science Foundation Ireland. His research focuses on distributed and adaptive resource management in wireless networks. Recent and current research sponsors include the National Science Foundation, the Science Foundation Ireland, the European Commission, and DARPA. He is an IEEE Communications Society Distinguished Lecturer from 2015 to 2018, for contributions to cognitive networking and to resource management in wireless networks.