

Dimensioning spectrum to support ultra-reliable low-latency communication

André Gomes*, Jacek Kibilda*, Nicola Marchetti†, and Luiz A. DaSilva*

*Commonwealth Cyber Initiative, Virginia Tech, USA, E-mail: {gomesa,jkibilda,ldasilva}@vt.edu

†CONNECT, Trinity College Dublin, Ireland, E-mail: {nicola.marchetti}@tcd.ie

Abstract—Industry-led initiatives such as the Next G Alliance (NGA) are currently considering how to dimension the spectrum required to support new classes of services envisioned beyond 5G. In particular, support for ultra-reliable low-latency communication (URLLC) brings the challenge of how to dimension stochastic wireless networks to meet stringent reliability and latency requirements. Our analysis indicates that the bandwidth needed to meet URLLC goals can be in the order of gigahertz, beyond what is available in today’s mobile networks. Network densification can ease those bandwidth needs but requires new deployment strategies involving substantially larger numbers of sites. As an alternative, we consider multi-connectivity and multi-operator network sharing as efficient ways to reduce the demand for bandwidth without outright deployment of additional base stations.

I. INTRODUCTION

Industry-led initiatives such as the Next G Alliance (NGA), an initiative to advance the next generations of wireless networks in North America, are currently considering how to dimension the spectrum to support new classes of services envisioned beyond 5G [1]. As communication services are identified, dimensioning spectrum corresponds to identifying suitable key performance indicators (KPIs) and assessing the amount of spectrum needed to meet the service-level requirements. This is key for future licensing, regulatory, and network deployment plans, and can even decide the fate of emerging communication services. In particular, ultra-reliable low-latency communication services pose a challenge of how to dimension the spectrum in stochastic wireless networks to meet stringent reliability and latency requirements. For instance, a typical URLLC requirement specified by the 3GPP is the transmission of 32 bytes within a 1 ms deadline with a 99.999% success probability, making URLLC sensitive to even rare events that occur with the probability of 0.001%.

A natural way to compensate for uncertainty in wireless communication, increase reliability, and reduce latency is to provision additional resources, be it in the form of additional antennas (e.g., multi-antenna communication in [2]), spectrum (e.g., wider channels in [2] and channel hopping in [3, §III.B.3]), base stations (e.g., network densification in [4]), or connections (e.g., multi-connectivity in [5], [6], [7]). Similarly, several works focus on the prioritization of URLLC traffic, which is also a form of provisioning additional resources, as more resources are made available to URLLC services compared to other types of services (e.g., scheduling prioritization of URLLC traffic over eMBB traffic in [8]).

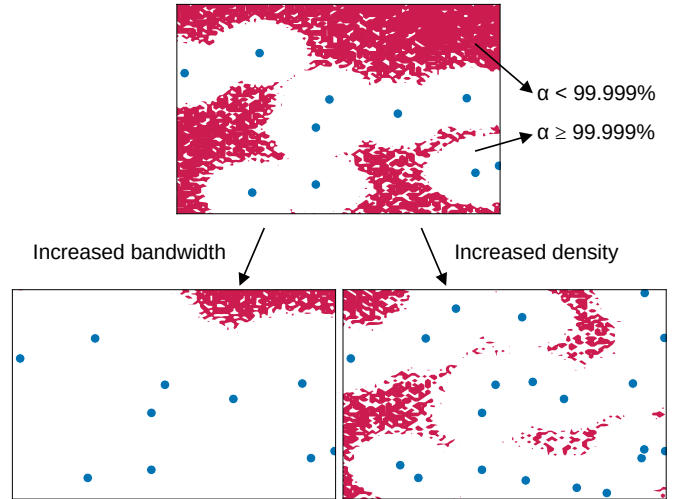


Fig. 1: The provision of additional bandwidth and/or network density expands the reach of URLLC (white zones, with reliability $\alpha \geq 99.999\%$). Blue markers depict base stations.

These, and other similar techniques discussed in [3], [4], show that supporting URLLC services requires that additional resources be made available. Yet, it remains an open question what level of resources is needed to achieve URLLC goals in large-scale mobile networks.

In this paper, we address this question and study the fundamental limits on the magnitude of resources needed to deploy large-scale mobile networks capable of supporting URLLC services. In particular, we seek to determine how to dimension the spectrum. The objective is to guide ongoing and future spectrum standardization efforts as well as future network deployment initiatives regarding the type and magnitude of spectrum required to enable URLLC services in mobile networks. To that end, we first map the URLLC service-level requirements onto bandwidth requirements and then assess the magnitude and type of spectrum needed to meet the service-level requirements on urban-macro network models developed by the 3GPP. As illustrated in Fig. 1, making more bandwidth available compensates for cell edge effects and expands the reach of URLLC. However, our results indicate that supporting URLLC services at large-scale networks can demand bandwidth in the order of gigahertz, beyond what is typically available in today’s networks. We study how network densification can expand the reach of URLLC services without

additional bandwidth, as illustrated in the right side of Fig. 1. We then study multi-connectivity and multi-operator network sharing as alternative ways to further ease the demand for bandwidth without outright deployment of additional base stations.

The remainder of this paper is organized as follows. § II describes the mapping between URLLC requirements onto bandwidth, followed by a description of the network models considered in this paper. § III presents the numerical results on the magnitude and type of spectrum needed to support URLLC services according to the network models developed by the 3GPP. Finally, § III summarizes the main takeaways from this paper.

II. PRELIMINARIES

A. Expressing reliability

We adopt the ITU-R/3GPP definition of reliability for URLLC [9], [10]: the success probability α of transmitting δ bits of data within a user plane latency deadline τ , i.e., $\Pr(T \leq \tau) \geq \alpha$. We consider a typical mobile-to-base station link, without intermediate relays. The user plane latency T , i.e., the time it takes to transmit δ bits from/to the base station to/from the mobile in the downlink/uplink communication [10], is a sum of physical and data link layers' delays, which generally includes (a) a processing delay, incurred by signal processing; (b) a medium access delay, incurred by scheduling or contention for the medium; and (c) a transmission delay, i.e., $T = T_{\text{proc}} + T_{\text{ma}} + T_{\text{tx}}$. The processing delay is fairly constant and hardware/implementation related; the medium access delay can be deterministic (e.g., scheduling-based approaches) or random (e.g., contention-based approaches). In URLLC, deterministic medium access is likely to be adopted, given the stringent reliability and latency requirements.

Hence, the transmission delay is the prime source of uncertainty given the probabilistic nature of wireless communication, prone to stochastic channel conditions caused by user mobility, blockages, fading, and interference. We adopt a general formulation where the transmission delay relates to the Shannon-Hartley capacity and assume no protocol overhead. This way, we can estimate the *fundamental minimum amount of spectrum* needed to support URLLC services. The transmission delay is expressed as:

$$T_{\delta} = \frac{\delta}{w \times \log_2(1 + \Gamma)}, \quad (1)$$

where w and Γ are the bandwidth and signal-to-interference-plus-noise ratio (SINR), respectively. We assume the SINR to be static during time τ , which is a reasonable assumption for small values of τ , such as in URLLC, where τ is often 1 ms.

With the focus on the transmission delay, the URLLC goals can be expressed as:

$$\Pr\left(w \geq \frac{\delta}{\tau_{\text{tx}} \times \log_2(1 + \Gamma)}\right) \geq \alpha, \quad (2)$$

where τ_{tx} is the fraction of the latency budget τ reserved for the transmission delay. Hence, the minimum bandwidth to satisfy

the URLLC service-level requirements is a function of the α -quantile Q_{α} of the inverse of the spectrum efficiency,

$$w = \frac{\delta}{\tau_{\text{tx}}} \times Q_{\alpha}\left(\frac{1}{\log_2(1 + \Gamma)}\right). \quad (3)$$

B. The network model

We are interested in the magnitude of spectrum needed to support URLLC in a mobile network. To that end, we adopt the 3GPP urban-macro network model in TR 38.901 [11]. The model defines (a) antenna array configurations for base stations [11, §7.3]; (b) path loss [11, §7.4]; (c) fading [11, §7.4]; and (d) line-of-sight probabilities [11, §7.4] for different frequency bands ranging from 0.5 to 100 GHz. We consider three frequency bands centered at 700 MHz, 4 GHz, and 30 GHz to study low, mid, and high bands (as recommended by the 3GPP in TR 38.913 [10]) and base station antenna arrays of size 2x2, 4x4, and 8x8, respectively. Further parameters, in compliance with the aforementioned reports, are as follows: (a) base stations: i. 25 m high and ii. 49 dBm transmit power; (b) mobile user: i. 1.5 m high, ii. 0 dBi omnidirectional antenna, and iii. 9 dB noise figure; (c) -90 dBm noise floor.

III. WHAT ARE THE MINIMUM SPECTRUM REQUIREMENTS FOR THE MOBILE NETWORK TO MEET URLLC GOALS?

In this section, we study the magnitude of spectrum needed to deploy URLLC-enabled mobile networks. Let us consider the case of $\delta = 32$ bytes and $\tau = 1$ ms [10]. We assume processing and medium access delays to be negligible. In practice, processing and medium access delays consume a fraction of the delay budget τ ; however, these delays can be considered deterministic in URLLC, as discussed in § II-A. We focus on the transmission delay, for it involves the primary sources of uncertainty in wireless networks, meaning that the transmission budget takes the entire latency budget, i.e., $\tau_{\text{tx}} = \tau$ in Eq. (2). This allows us to quantify the *minimum* amount of spectrum needed to support URLLC services. We focus on three reliability regimes, which, from now on, we denote as: (a) reliable ($\alpha = 90\%$); (b) highly-reliable ($\alpha = 99.9\%$); and (c) ultra-reliable ($\alpha = 99.999\%$).

We consider two downlink mobile communication scenarios. The first scenario corresponds to mobile networks subject to stochastic path loss and fading, and negligible interference, which we refer to as a noise-limited network, reflecting networks with a low frequency reuse or tight coordination between transmitters. The second scenario, an interference-limited network, includes interference, representing mobile networks subject to full frequency reuse, with little or no interference coordination between transmitters. Practical networks may operate somewhere in between noise- and interference-limited models, for they may adopt some forms of interference mitigation. Nevertheless, we consider noise- and interference-limited networks as the *lower-* and *upper-bound* estimates of the *fundamental minimum* amount of spectrum necessary to enable URLLC services in mobile networks.

Our results stem from system-level Monte Carlo experiments. In each experiment, we consider the performance of

a typical mobile user placed at the origin and base stations deployed according to a Poisson point process (PPP) of density λ base stations (BSs) per km^2 . The typical mobile associates with the base station of the highest average received power. We assume perfect beam alignment between mobiles and serving base stations and a fixed network load of one active connection per base station. This way, interference at the typical mobile is caused by the beam alignment between non-serving base stations and other randomly placed mobiles in the network. We generate 3×10^7 experiments for each network density to capture even low-probability channel conditions in the order of 0.001% when $\alpha = 99.999\%$.

A. Noise-limited networks

Fig. 2 shows the demand for bandwidth to reach certain levels of reliability in different network densities for the three frequency bands of interest, 700 MHz, 4 GHz, and 30 GHz. Let us consider a mobile network of density 30 BS/ km^2 , a density analogous to dense urban networks in [10], at the 30 GHz band as a reference. This network can be tailored to support reliable (i.e., $\alpha = 90\%$) to ultra-reliable communication (i.e., $\alpha = 99.999\%$) by the provision of an order of magnitude additional bandwidth (see vertical red arrows in Fig. 2). Alternatively, increased density can ease the demand for bandwidth, such as illustrated by the horizontal red arrow in Fig. 2, where ultra-reliable communication demands as much as bandwidth as highly-reliable communication (i.e., $\alpha = 99.9\%$) at the cost of a roughly $2.3\times$ denser deployment.

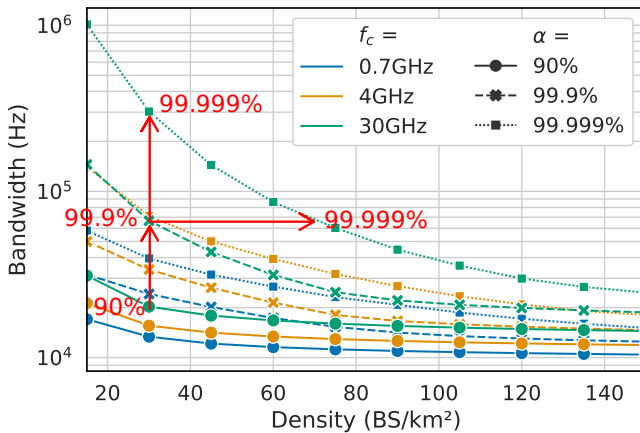


Fig. 2: The bandwidth required to meet the URLLC requirement of $\delta = 32$ bytes and $\tau = 1$ ms in noise-limited networks operating in different frequency bands.

The relationship between bandwidth and network density changes for different reliability levels α . An increase in network density can be traded for up to an order of magnitude of bandwidth in ultra-reliable communication (dotted lines), whereas it has marginal impact on less reliable scenarios, such as $\alpha = 90\%$ (continuous lines). This is a consequence of how network density impacts coverage. As illustrated in Fig. 1, the farther the mobile is from the base station, the harder it is to meet the reliability requirements because of weak coverage, increasing the demand for additional bandwidth. Increased

network density improves coverage by reducing the distance between transmitter and receiver, as well as the probability of non-line-of-sight communication. As ultra-reliable communication is more sensitive to long distance communication than less reliable regimes, ultra-reliable communication benefits the most from increased network density.

B. Interference-limited networks

Interference-limited networks demand additional spectrum to compensate for interference. Fig. 3 is Fig. 2's counterpart and shows that achieving highly- (i.e., $\alpha = 99.9\%$) and ultra-reliable communication (i.e., $\alpha = 99.999\%$) can demand orders of magnitude more bandwidth than in noise-limited networks, whereas the requirements are somewhat similar for reliable communication (i.e., $\alpha = 90\%$) in both networks. Our system model assumes highly directional antennas with perfect beam alignment, which naturally decreases the overall interference, being it marginal when $\alpha \leq 90\%$. However, as the reliability requirement increases, communication is more sensitive to even low-probability interference conditions, such as in edge-case scenarios when nearby non-serving base stations (i.e., interferers) steer their beams towards the mobile, demanding wider bandwidth to compensate for strong interference.

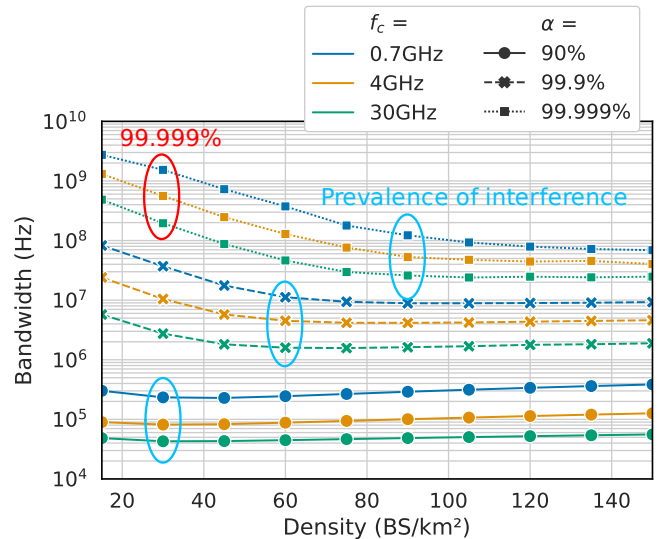


Fig. 3: The bandwidth required to meet the URLLC requirement of $\delta = 32$ bytes and $\tau = 1$ ms in interference-limited networks operating in different frequency bands. The red circle indicates the bandwidth required for ultra-reliable communication when the network density is 30 BS/ km^2 . Blue circles indicate the network density where interference prevails for different reliability levels.

The amount of bandwidth to support ultra-reliable communication is beyond what is typically available in today's networks. In dense urban networks, where network density is approximately 30 BS/ km^2 [10], the required bandwidth to meet ultra-reliable communication is around a gigahertz, which is impossible in the 700 MHz band and impractical in the 4 GHz band. The substantial amount of bandwidth favors

the use of higher frequency bands, where wider bandwidth is typically available. Furthermore, higher frequency bands are often coupled with denser antenna arrays (as in our network model), enabling higher directionality transmission that reduces interference and increases resilience to path loss. This translates into less spectrum required to support URLLC services than in the case of lower frequency bands.

Increased network density can ease the bandwidth requirements to some extent. Network density increases coverage until interference prevails, that is, dominant interfering links shift from non-line-of-sight to line-of-sight, and network densification starts to have a negligible (or even harmful) effect on coverage [12]. Interestingly, interference prevails at different network densities for different reliability levels. For instance, Fig. 3 suggests that increased network density leads to marginal returns in reliable communication, where interference prevails at approximately 30 BS/km². On the other hand, densification substantially eases the demand for bandwidth in highly- and ultra-reliable communication up to the prevalence of interference at approximately 60 BS/km² and 90 BS/km², respectively.

Another way to look at the impact of network density is to consider the delay at a fixed bandwidth. Fig. 4 illustrates the mapping from α and T_{tx} to network density when bandwidth is fixed at 100 MHz in the 30 GHz band, the same amount of bandwidth recently auctioned in the US in auction 101 of the Federal Communications Commission (FCC) for one block of spectrum in upper 30 GHz bands. Lower reliability regimes (e.g., $\alpha \leq 99\%$) are only marginally impacted by network density; in fact, densification tends to increase the transmission delay in that region, as shown in the upper left side Fig. 4. On the other hand, as we increase the reliability requirements (e.g., $\alpha \geq 99.999\%$), there is a substantial demand for increased network density. For instance, a general URLLC requirement of $\alpha = 99.999\%$ and $\tau = 1$ ms demands a network density of 60 BS/km² or denser in Fig. 4.

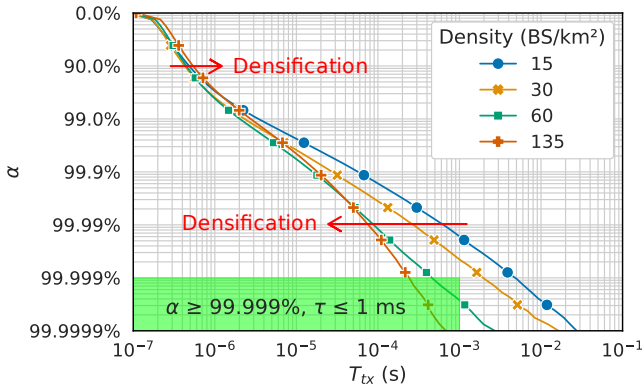


Fig. 4: The impact of network density on the transmission delay when bandwidth is fixed at 100 MHz in the 30 GHz band. A general URLLC requirement of $\alpha = 99.999\%$ and $\tau = 1$ ms (green zone) demands densities ≥ 60 BS/km². Similar patterns were found for the 700 MHz and 4 GHz bands.

The distinct trends in network density for different re-

liability regimes are due to how network density impacts overall and edge coverage. Unlike in noise-limited networks, where network density directly maps onto signal strength, the relationship between network density and coverage is subtle in interference-limited networks, for increased network density increases both signal strength and interference. As shown in Fig. 5, increased network density incurs a small penalty for the overall coverage (e.g., coverage probability $\geq 10^{-2}$, or 1%); on the other hand, increased network density significantly increases edge coverage (e.g., coverage probability $< 10^{-2}$). (Results consistent with this finding are reported in [13]). The sensitivity to highly improbable events, such as edge conditions, increases with the reliability degree α (e.g., events with probability as low as 0.001% when $\alpha = 99.999\%$). As such, ultra-reliable communication is more sensitive to edge coverage than less reliable regimes, benefiting the most from increased network density.

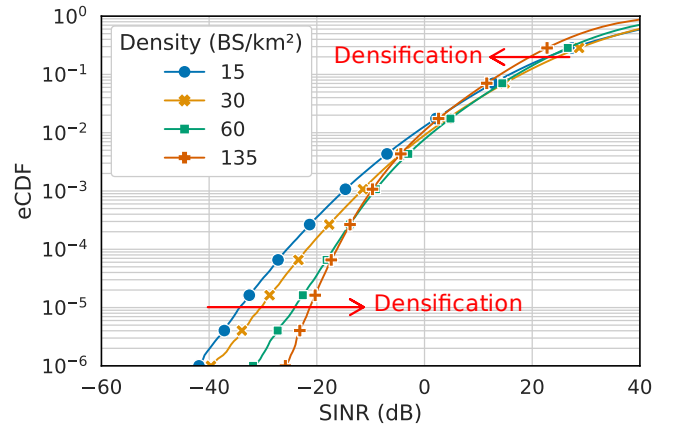


Fig. 5: The empirical cumulative distribution function of the SINR for different densities in the 30 GHz band. Similar patterns were found for the 700 MHz and 4 GHz bands.

C. Easing the bandwidth requirements

As discussed in the previous section, deploying mobile networks capable of supporting ultra-reliable communication ($\alpha \geq 99.999\%$) can demand massive amounts of resources (e.g., ≥ 1 GHz bandwidth or ≈ 100 BS/km²). In this section, we consider multi-connectivity and multi-operator connectivity as strategies to reduce this demand.

Multi-connectivity is an enabler of URLLC that leverages diversity and redundancy in the form of multiple connections to mitigate the volatility of a wireless link [3], [6], [7]. In multi-connectivity, a mobile multi-connects to distinct base stations. The received signals can be combined in several ways using diversity combining techniques. In our analysis, we assume signals are weighted and summed according to the respective signal strengths, as in maximum ratio combining.

Multi-operator connectivity combines multi-connectivity and multi-operator network sharing. In multi-operator connectivity, the mobile can multi-connect to base stations of different mobile operators in a multi-operator network. The multi-operator network sharing arrangement is akin to mobile virtual

network operators like Google Fi, in the US, which operates atop several mobile operators. However, unlike Google Fi-like multi-operator networks, where users are often subject to the same radio resource management policy as subscribers of the underlying operators, we assume that bandwidth can be reserved and isolated for URLLC users at each operator, for instance through network slicing. Furthermore, we assume users can simultaneously connect to base stations of more than one operators, as opposed to a single connection to a single base station at a time.

The resources in multi-operator networks differ from traditional single-operator networks in two ways. First, the base stations exhibit spatial correlation because of clustering between the involved operators [14]. Intuitively, clustering is caused by similar network planning strategies by the operators, such as covering areas of mutual interest, such as dense central business districts, hence deploying base stations adjacent to each other, or sharing towers, hence co-locating infrastructure. Second, multi-operator connectivity accounts for greater spectrum diversity because each mobile operator operates in its own licensed spectrum.

As in § III-B, we consider interference-limited networks. Fig. 6 shows the bandwidth required to meet the URLLC requirement of $\delta = 32$ bytes, $\tau = 1$ ms, and $\alpha = 99.999\%$. We focus our analysis on the 30 GHz frequency band; nevertheless, we can report that similar patterns were found in the 700 MHz and 4 GHz bands. Here, single-connectivity consists of a single connection to a single operator (as in previous sections; hence our baseline); multi-connectivity, two connections, each to a base station of the same operator; and multi-operator connectivity, two connections, one connection to a base station of operator A and one connection to a base station of operator B. We model clustering in multi-operator connectivity as a Gauss-Poisson point process where clusters are deployed according to a PPP of density $\lambda/2$ [15]. Each cluster consists of two points: the first is at the center of the cluster and corresponds to a base station of operator A, and the second is randomly placed within a radius of 50 m and corresponds to a base station of operator B. Each mobile operator operates in its own licensed band with full frequency reuse.

Multi-connectivity reduces the required bandwidth by approximately one order of magnitude. The reduction in the amount of bandwidth stems from the spatial diversity introduced by simultaneous connections to two base stations, which results in connections with distinct line-of-sight, path loss, and fading patterns. On the downside, both connections are subject to the same interference component because of full frequency reuse by the operator.

Multi-operator connectivity, in turn, benefits from frequency diversity, as each operator operates in its own licensed spectrum. As a result, multi-operator connectivity further reduces the required bandwidth, approximately two orders of magnitude compared to the baseline, from hundreds of megahertz to a few megahertz when the network density is 30 BS/km². Clustering incurs a small penalty in multi-operator connectivity compared to what it would be if the operators were unclustered (i.e., independent deployments), marginally increas-

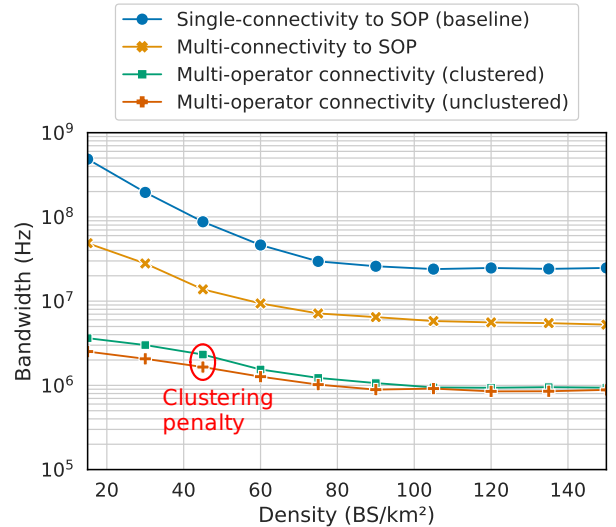


Fig. 6: The bandwidth required to meet the URLLC requirement of $\delta = 32$ bytes, $\tau = 1$ ms, and $\alpha = 99.999\%$ at the 30 GHz band in interference-limited networks. SOP stands for single-operator networks. The multi-operator network in multi-operator connectivity consists of two operators.

ing the demand for bandwidth. Furthermore, multi-operator connectivity is somewhat insensitive to the network density compared to single-connectivity and multi-connectivity and does not require massive network densification to reduce the demand for bandwidth, hence potentially reducing the cost of deploying networks capable of supporting URLLC services.

IV. CONCLUSION

Dimensioning spectrum is key to ongoing and future spectrum standardization efforts as well as future network deployment campaigns to enable URLLC services. In this paper, we studied how to dimension the spectrum for URLLC services by first mapping the URLLC requirements onto bandwidth and then assessing the amount of bandwidth it takes to support URLLC services in large-scale 3GPP mobile networks. Our study indicates that the demand for bandwidth heavily depends on the reliability requirement of interest. In ultra-reliable communication, where reliability is $\geq 99.999\%$, the amount of bandwidth can be in the order gigahertz, as opposed to less than a megahertz when reliability is 90%. Higher frequency bands, often coupled with wider bandwidth and denser antenna arrays than lower bands, are required in order to meet ultra-reliability goals. The demand for bandwidth can be traded for network density, but it requires massive densification (e.g., ≈ 100 BS/km²) to significantly lower the required bandwidth, which is costly or even impractical, depending on the spatial availability for deployment of base stations. Alternatively, multi-operator connectivity, a form of multi-connectivity to multiple mobile operators, was shown to reduce the demand for bandwidth by a few orders of magnitude without the outright deployment of a massive number of base stations.

As future work, we plan to investigate other sources of randomness in the network, such as network load, and ways to

circumvent their impact on the demand for network resources in the context of URLLC, such as using traffic prediction techniques. As a byproduct of this paper, we also plan to study the problem of efficiently estimating stringent reliability KPIs. As in this paper, estimating the bandwidth at the 99.999% reliability level is time-consuming because it takes $> 10^5$ data samples (or simulation cycles) to capture rare network conditions that may occur only 0.001% of the time, which can be impractical depending on the time it takes to collect each sample and the time window available for data collection and assessment.

ACKNOWLEDGMENTS

The research leading to this paper received support from the Commonwealth Cyber Initiative (CCI) in Virginia, the US, an investment in the advancement of cyber R&D, innovation, and workforce development: www.cyberinitiative.org. It was also supported by the Science Foundation Ireland under grants 17/NSFC/5224 and 13/RC/2077_P2 (CONNECT).

REFERENCES

- [1] Next G Alliance, "Next G alliance report: roadmap to 6G," Next G Alliance, Tech. Rep., 2 2022.
- [2] O. N. Yilmaz, Y.-P. E. Wang, N. A. Johansson, N. Brahma, S. A. Ashraf, and J. Sachs, "Analysis of ultra-reliable and low-latency 5G communication for a factory automation use case," in *2015 IEEE international conference on communication workshop (ICCW)*. IEEE, 2015, pp. 1190–1195.
- [3] G. J. Sutton, J. Zeng, R. P. Liu, W. Ni, D. N. Nguyen, B. A. Jayawickrama, X. Huang, M. Abolhasan, Z. Zhang, E. Dutkiewicz, and T. Lv, "Enabling technologies for ultra-reliable and low latency communications: From PHY and MAC layer perspectives," *IEEE Communications Surveys Tutorials*, vol. 21, no. 3, pp. 2488–2524, 2019.
- [4] P. Popovski, J. J. Nielsen, C. Stefanovic, E. De Carvalho, E. Strom, K. F. Trillingsgaard, A.-S. Bana, D. M. Kim, R. Kotaba, J. Park *et al.*, "Wireless access for ultra-reliable low-latency communication: principles and building blocks," *IEEE Netw.*, vol. 32, no. 2, pp. 16–23, 2018.
- [5] A. Gomes, J. Kibilda, A. Farhang, R. Farrell, and L. A. Da Silva, "Multi-operator connectivity sharing for reliable networks: A data-driven risk analysis," *IEEE Transactions on Network and Service Management*, pp. 1–1, 2021.
- [6] A. Wolf, P. Schulz, M. Dörpinghaus, J. C. S. Santos Filho, and G. Fettweis, "How reliable and capable is multi-connectivity?" *IEEE Trans. Commun.*, vol. 67, no. 2, pp. 1506–1520, 2018.
- [7] 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, 2018.
- [8] A. Anand, G. De Veciana, and S. Shakkottai, "Joint scheduling of URLLC and eMBB traffic in 5G wireless networks," *IEEE/ACM Transactions on Networking*, vol. 28, no. 2, pp. 477–490, 2020.
- [9] ITU-R, "Minimum requirements related to technical performance for IMT-2020 radio interface(s)," International Telecommunication Union (ITU), Tech. Rep. ITU-R M.2410-0, 11 2017.
- [10] 3GPP, "Study on scenarios and requirements for next generation access technologies," 3rd Generation Partnership Project (3GPP), Tech. Rep. 38.913, 07 2020, version 16.0.0.
- [11] —, "Technical specification group radio access network; study on channel model for frequencies from 0.5 to 100 GHz," 3rd Generation Partnership Project (3GPP), Tech. Rep. 38.901, 12 2019, version 16.1.0.
- [12] M. Ding, P. Wang, D. López-Pérez, G. Mao, and Z. Lin, "Performance impact of LOS and NLOS transmissions in dense cellular networks," *IEEE Transactions on Wireless Communications*, vol. 15, no. 3, pp. 2365–2380, 2016.
- [13] M. Rebato, M. Mezzavilla, S. Rangan, F. Boccardi, and M. Zorzi, "Understanding noise and interference regimes in 5G millimeter-wave cellular networks," in *European Wireless 2016; 22th European Wireless Conference*, 2016, pp. 1–5.
- [14] J. Kibilda, B. Galkin, and L. A. DaSilva, "Modelling multi-operator base station deployment patterns in cellular networks," *IEEE Transactions on Mobile Computing*, vol. 15, no. 12, pp. 3087–3099, 2016.
- [15] M. Haenggi, *Stochastic geometry for wireless networks*. Cambridge University Press, 2012.

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

Jacek Kibilda (M'16, SM'19) is the 5G and AI Research Assistant Professor with the Commonwealth Cyber Initiative and a Research Assistant Professor with the Bradley Department of Electrical and Computer Engineering at Virginia Tech. Previously he was a Research Fellow with CONNECT at Trinity College Dublin and a Challenge Research Fellow with Science Foundation Ireland. He is a Fulbright fellow class 2017-18. Jacek's research focuses on modeling and technology design for Next G mobile networks.

Nicola Marchetti (Senior Member, IEEE) received the Ph.D. degree in wireless communications from Aalborg University, Denmark, and the M.Sc. degrees in electronic engineering and mathematics from the University of Ferrara, Italy, and Aalborg University. He is currently an associate professor in wireless communications with Trinity College Dublin, Ireland. He performs his research under CONNECT, where he leads the Wireless Engineering and Complexity Science (WhyCOM) Lab. His research interests include Autonomous and Self-Organising Networks, Communications for Biology, PHY Algorithms and MAC Protocols, and Quantum Communications and Networks. He has authored in excess of 160 journals and conference papers, and received four best paper awards. He has served as an associate editor of the IEEE Internet of Things Journal and the EURASIP Journal on Wireless Communications and Networks, and is a Fellow of Trinity College Dublin.

Luiz A. DaSilva is the Executive Director of the Commonwealth Cyber Initiative (CCI), and the Bradley Professor of Cybersecurity at Virginia Tech. He previously held the chair of Telecommunications at Trinity College Dublin, where he served as the Director of CONNECT, the telecommunications research centre funded by the Science Foundation Ireland. Prof. DaSilva's 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. Prof. DaSilva is a Fellow of the IEEE, for contributions to cognitive networking and to resource management in wireless networks, and has been an IEEE Communications Society Distinguished Lecturer (2015-18) and a Fellow of Trinity College Dublin.