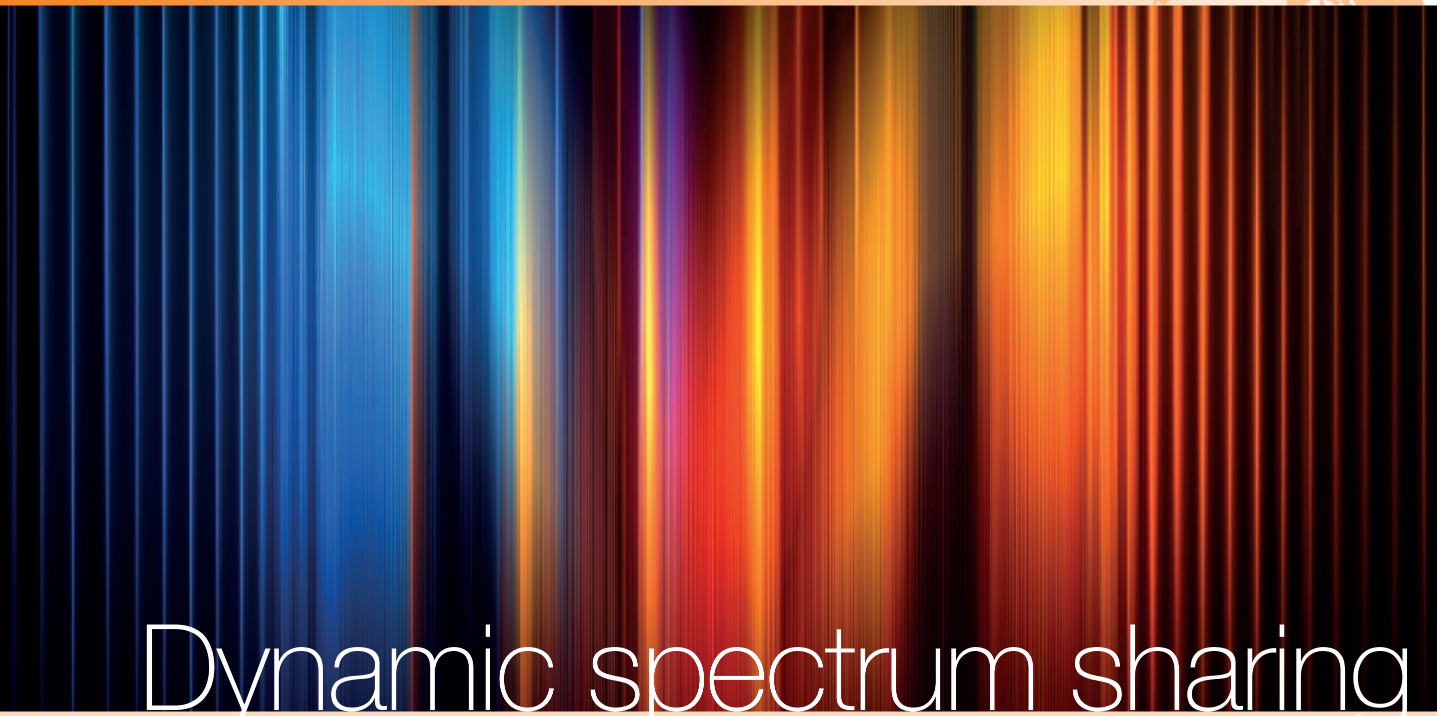

Dynamic spectrum sharing

between active and passive users above 100GHz

by Michele Polese, Viduneth Ariyarathna, Priyangshu Sen, Jose V. Siles, Francesco Restuccia, Tommaso Melodia & Josep M. Jornet

Sixth-generation wireless networks will aggregate higher-than-ever mobile traffic into ultra-high capacity backhaul links, which could be deployed on the largely untapped spectrum above 100GHz. Current regulations however prevent the allocation of large contiguous bands for communications at these frequencies, since several narrow bands are reserved to protect passive sensing services. These include radio astronomy and Earth exploration satellites using sensors that suffer from harmful interference from active transmitters. Here we show that active and passive spectrum sharing above 100GHz is feasible by introducing and experimentally evaluating a real-time, dual-band backhaul prototype that tracks the presence of passive users (in this case the NASA satellite Aura) and avoids interference by automatically switching bands (123.5–140GHz and 210–225GHz). Our system enables wide-band transmissions in the above-100GHz spectrum, while avoiding harmful interference to satellite systems, paving the way for innovative spectrum policy and technologies in these crucial bands.

The digital transformation of our society is fostered by the availability of a fundamental, invisible, yet scarce resource – the electromagnetic spectrum. Besides enabling information exchange through wireless communications, the electromagnetic spectrum is also a rich source of information through sensing. The finite nature of the spectrum creates competing interests for communications and sensing. These diverging interests, expressed by different scientific communities, government entities and industries, have led to rigid spectrum allocations from national and international regulatory bodies, such as the International Telecommunications Union (ITU), or the Federal Communications Commission (FCC), dating back to the 1930s.

To support more devices and ultra-high capacity applications, 6th Generation (6G) wireless networks will require data rates that are orders of magnitude higher than available today, thus boosting the need for spectrum. While the 5th Generation (5G) of mobile networks uses carrier frequencies as high as 71GHz, 6G will move beyond 100GHz to aggregate data of many mobile users in ultra-high capacity backhaul links.

However, communications in this spectrum band are limited by the coexistence of passive users that (i) do not transmit and (ii) only use high-sensitivity Radio Frequency (RF) sensors for Earth exploration, weather monitoring and radio-astronomy. Passive users can be negatively impacted by interference from active transmissions. Therefore, they retain exclusive access to relatively narrow portions of the above-100GHz spectrum, pre-venting the allocation of contiguous chunks with tens of GHz of bandwidth for communications. For example, in the U.S., the largest allocations for active transmissions between

100 and 275GHz are 32.5GHz (116–148.5GHz) and 18.5GHz (231.5–250GHz), but with only 12.25GHz (non contiguous) earmarked for unrestricted fixed or mobile terrestrial usage[6]. Transmissions are strictly forbidden in 33.5GHz of spectrum, and conditioned to the protection of coexisting passive users in the remaining spectrum.

These conservative regulations apply even without passive users exploiting the spectrum for sensing. This prevents multiplexing of unused resources, making this spectrum less appealing for wireless backhaul. To foster wireless innovation in years to come, it becomes necessary to develop spectrum sharing solutions between communications and passive sensing systems – which is the main contribution of this paper. Moreover, while terrestrial sensing stations can be protected through geographical separation, orbiting satellite systems call for dynamic sharing solutions based on self-adaptive wireless links. Indeed, we show—through an accurate link budget based on ITU channel models—that active transmitters may actually generate harmful interference to highly sensitive sensors on satellites orbiting over communication systems.

Shared access between sensing and communications has never been demonstrated experimentally in the spectrum above 100GHz. This is not without reason. Current spectrum policies prohibit transmissions in passive bands, even for experimental purposes. Additionally, a self-adaptive wireless link requires the design of dynamic sub-terahertz RF front-ends and protocol stacks that need to be controlled and reconfigured in real time. However, existing transceivers designed for this portion of the spectrum are limited to channel sounding or static baseband processing capabilities, with limited or no reconfigurability, and are thus unfit to demonstrate dynamic spectrum access.

This article presents a dynamic spectrum access, dual-band backhaul system in the above-100GHz spectrum, specifically, at 123.5–140GHz (referred to as Lower Band (LB) throughout the paper) and 210–225GHz (Upper Band (UB)). This is the first above-100GHz wireless system that (i) operates a real-time protocol stack in two frequency bands; (ii) performs automated and dynamic spectrum sharing; and (iii) adapts link parameters (e.g., the frequency band) in real time. The proposed system tracks satellite orbits and automatically switches frequency bands, thus preventing backhaul transmissions from interfering with passive incumbents, while enabling the allocation of large, contiguous bands for communication services on shared spectrum. Prototyping this system has required us to address challenges that span from RF design for the dual-band setup, to dynamic control and integration of multiple independent systems through open interfaces. We evaluate the backhaul prototype in an experimental setting with transmitter and receiver deployed outdoor in an urban environment, on the roofs of two buildings, under the limits of an FCC experimental license. We believe that the results shown in this paper can pave the way for more flexible regulations, protecting scientific uses of the spectrum while allowing more reasonable allocations for future generation of communication systems.

Results

This section introduces the proposed self-adaptive dual-band backhaul link. We discuss the challenges and lessons learned in the development of the RF and communication components, the dynamic switching framework, and present its experimental performance evaluation.

RF and communications components

Figure 1 shows the prototype hardware, with Transmitter (TX)/Receiver (RX) real-time signal processing back-ends and dual-band sub-terahertz front-ends. The NI (formerly known as National Instruments) mmWave back-end features a hardware-accelerated protocol stack that generates an Orthogonal Frequency Division Multi-plexing (OFDM) waveform inspired by the 5G New Radio (NR) physical (PHY) layer, aggregating eight component carriers with 100MHz each. The dual-band front-ends, supporting 123.5–140GHz and 210–225GHz, accept an Intermediate Frequency (IF) input at the up-converters (TX-side) and generate an IF output at the down-converters (RX-side). They are connected to the baseband-to-IF conversion modules of the NI system through a Single-Pole Double-Throw (SPDT) RF switch, which supports DC–18GHz. Four Keysight E8257D Performance Signal Generators (PSGs) drive the Local Oscillators (LOs) of the dual-band front-ends.

The overall system design was hindered by two main challenges. First, the same IF output, with high peak-to-average power ratio, needs to be interfaced with dual-band front-ends with different and limited input power specifications. Second, phase noise and dual sideband front-ends introduce amplitude impairments in the received signal.

The link budget planning aims at having a 10 decibel (dB) SNR margin at the IF input of the receiver back-end, independently on the LB or UB front-ends. This is challenged by two factors.

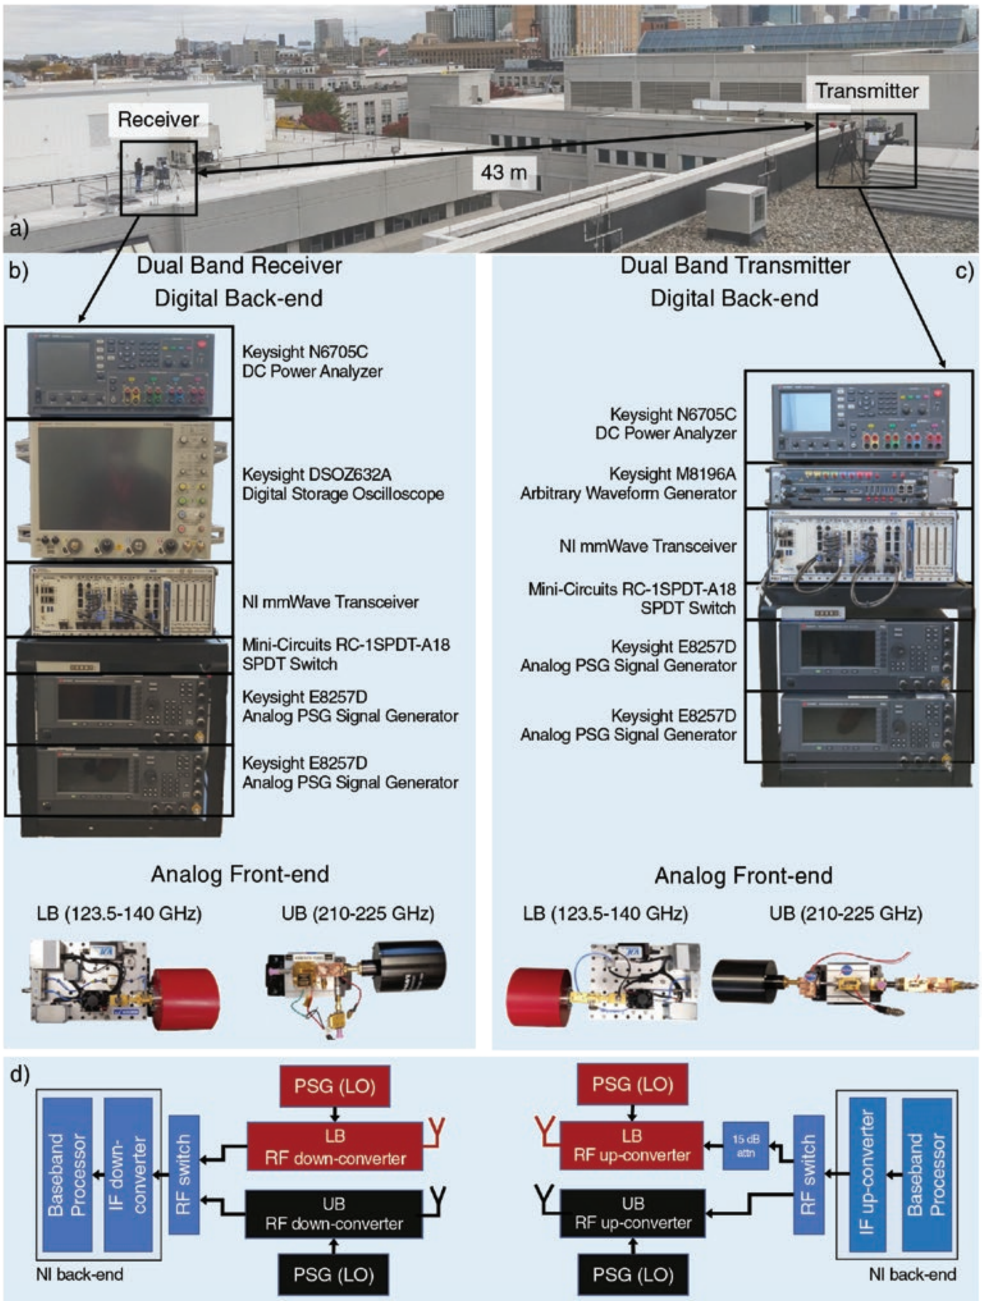

Fig. 1 Prototype schematics and dual-band backhaul link deployment.

a Deployment of the dual-band backhaul link over the Snell Engineering and Egan Research buildings at Northeastern University.

b Components of the RX.

c Components of the TX.

d Block diagram with the components of the TX and RX.

DC stands for Direct Current, NI for National Instruments, SPDT for Single-Pole Double-Throw, PSG for Performance Signal Generator, LB for Lower Band (123.5–140 GHz), UB for Upper Band (210–225 GHz), LO for Local Oscillator, IF for Intermediate Frequency, and RF for Radio Frequency.

First, the back-end OFDM implementation introduces a high peak-to-average power ratio (>10 dB), even for the lowest order modulations. This is a known drawback for OFDM systems using large Fast Fourier Transform (FFT) sizes (2048 for each component carrier, in this case). Second, there is an input power mismatch between the LB and UB TX front-ends. The LB front-end needs an input signal with power lower than 0 decibel-milliwatts (dBm), and generates a maximum 13dBm output power. The UB front-end needs an input power of 10dBm to generate a 3dBm output at the desired harmonic. To address this, we precisely characterised the OFDM waveform power, and reduced the back-end output power to maintain the peak power under the limit of the UB front-end. The input to the LB front-end was then further attenuated to be below 0dB. Finally, to close the link margin for the 43m link, the LB and UB systems use 38dBi and 40dBi antennas, respectively.

The TX and RX at both bands feature two mixing stages, part of the analogue front-ends shown in Fig. 1, thus forming a heterodyne system. However, due to the lack of RF filters at the considered frequency ranges, both LB and UB front-ends lack a sideband selection filter. This, combined with the high phase noise (due to the frequency multiplication process), may lead to drastic amplitude fluctuations in the received signal. To address this, we offset the IF frequency in the up- and down-converting paths by 1GHz, different from traditional heterodyne systems, which use the same IF frequencies. Thus, only one of the sidebands is selected for receiver-side processing, preventing undesired amplitude fluctuations. To enable quick dynamic switching, the LB and UB front-ends are set to the same IF frequency.

The frequency offset is compensated by appropriately setting the IF-to-baseband and the IF-to-RF LO frequencies in the NI back-end and the two LB and UB transmit-receive paths. Notably, the LB front-ends perform a 4 × frequency multiplication of the LO input, thus the TX LO is set to 32.5GHz (130GHz at the up- converter), while the RX LO is set to 32.75GHz (131GHz at the down-converter). Similarly, UB converters utilise an 9× (6×) LO multiplication at the TX (RX) side. Therefore, the LO input at the TX-side was set to 23.88GHz (corresponding to 215GHz) and RX-side was set to 216/6 = 36GHz, selecting the upper sideband for baseband processing.

Dynamic frequency switching framework

The dynamic control of the dual-band front-end requires the integration of the hard- ware components with a programmable, software-based control logic that can automatically switch bands when the active system interferes with a passive incumbent. The framework features (i) a tracking system that identifies when LB or UB transmissions may be interfering with a passive user (e.g., a satellite orbiting over the link); and (ii) an interface with the SPDT switch.

This flexible architecture can operate in two different modes. The first one is named independent switching, and assumes no coordination between the backhaul endpoints, which are only loosely time-synchronised to a common source through the Network Time Protocol (NTP) and perform frequency switching decisions independently.

This configuration represents an ad hoc network deployment, without out-of-band control. The framework runs on the NI back-ends. We call the second deployment mode coordinated switching. It mimics the configuration of backhaul nodes for cellular networks, where the equipment is connected to a control overlay. The tracking framework is hosted by a central controller, which notifies the two endpoints simultaneously. We considered as passive user the NASA satellite Aura, which studies the Earth climate and air quality and is listed as incumbent in the spectrum adjacent to the UB band in the World Meteorological Organization (WMO) Observing Systems Capability Analysis and Review Tool (OSCAR) database. Note that we operate the UB backhaul link in the 210–225GHz range, compliant with FCC regulations and an experimental license. However, its RF front-end could extend up to 240GHz, thus potentially interfering with the Aura Microwave Limb Scanner (MLS). This instruments scans the ozone (O3) absorption line around 235.7098GHz, and the carbon monoxide (CO) line at 230.538GHz. Therefore, in this specific setup, we consider LB as the band where UB traffic can be offloaded to avoid interference to the Aura MLS.

The tracking system retrieves Aura's orbits from public Application Programming Interfaces (APIs)[37] and checks whether the satellite is orbiting over the backhaul link deployment. Harmful interference is detected when the satellite and backhaul are in Line-of-Sight (LOS), or the actual link budget exceeds an ITU threshold (as discussed next). Accordingly, the SPDT switch interface triggers an UB-to-LB band switch through HyperText Transfer Protocol (HTTP) APIs. The switching mechanism can be further refined to avoid unnecessary or too frequent band switching that may degrade the system performance.

Experimental evaluation

The dual-band transmitter and receiver were deployed outdoor, on top of two Northeastern University campus buildings, at a distance of 43m, as shown in Fig. 1. Despite the similar link budget, the higher thermal and Single Sideband (SSB) phase noise (Fig. 2a, b) of the UB front-end (with a ~ 10dB noise figure increase) affects its throughput when higher Modulation and Coding Schemes (MCSs) are considered (Fig. 2c). The maximum throughput (achieved with the LB system) and the average throughput (which accounts for both) are comparable up to a Quadrature Phase Shift Keying (QPSK) modulation with a 1/2 coding rate. Less robust MCSs lead to lower throughput for UB. As a consequence, the average throughput is below the maximum achievable by the system.

To evaluate the effectiveness of the proposed frequency switching approach, Fig. 2d reports the throughput over time with independent and coordinated switching artificially triggered every 30s. The system maintains stable throughput with negligible impact of the switching procedure, especially with coordinated switching. This is only an illustrative example to show the effectiveness of the switching mechanism, as the actual time spent in each frequency band varies and depends on the presence of passive users in the UB band. Figure 2f plots the cumulative distribution function of the switching duration Δ_s over 120 independent switching events, defined as the time interval between the switch at one endpoint and the switch at the other. Independent switching requires a median value of Δ_s = 230 ms, as the switch may happen at slightly different time instants in the two endpoints, which only rely on NTP for a loose synchronisation. The coordinated approach takes <42 ms in the worst case (median 30.11 ms), as the switch is triggered by a centralised controller for both endpoints. A short switching time also allows the system to easily scale in case multiple incumbents need to be tracked and avoided. Overall, independent switching—which facilitates deployment in challenging environments—has a trade off in performance with coordinated switching. Even in case the frequency switch happens infrequently, and thus has limited impact on the average throughput, a downtime of tens of milliseconds may affect the end-to-end performance of traffic flows traversing the link, e.g., it may trigger retransmissions or session resets at the transport or application layers. Nonetheless, both switching mechanisms allow the prototype to promptly avoid harmful interference, while maintaining high average throughput values.

Discussion

To understand the dynamic spectrum sharing opportunities above 100GHz (specifically, between 209–241GHz), Fig. 3 illustrates spectrum regulations (bottom), and a link budget for the worst-case interference toward a sensing system on the Aura satellite orbiting over the backhaul link (top). For this analysis, we focus on the 209–241GHz band, as it is spectrum in which the UB front-end operates, but similar coexistence issues apply to the ITU-regulated spectrum above 100GHz.

The 209–241GHz spectrum includes bands where active users (i) cannot transmit, (226–231.5GHz, according to footnotes FCC US2466 and ITU 5.3405); or (ii) need to protect passive radio-astronomy and Earth monitoring users (209–226GHz and 235–238GHz, according to FCC US342, ITU 5.149, and 5.563). The ITU defines harmful interference as received power above the threshold $P_{RX,th}$ = – 194dBW in the 226 – 231.5GHz band. The top left corner of Fig. 3 evaluates whether this threshold is exceeded for Aura—which orbits 705km over the Earth—by reporting the received power P_{RX} for different elevation angles α. We conservatively assume that the backhaul link directional antenna (with a 40 dB gain, and 23dBm of transmit power) steers toward the satellite in a LOS scenario. This is usually not the case, yet provides a worst-case estimate of the interference. We consider the ITU channel models for the atmospheric attenuation as a function of frequency, height and elevation angle α. The carrier frequency is set to fc = 226 GHz, i.e., the closest frequency to the UB carrier actually used in the experiment (224GHz) with a threshold defined in ref. 40. While the atmospheric attenuation of the signal prevents interference for small α, for $\alpha \geq 8.4°$ the received power exceeds $P_{RX,th}$. This translates into satellite passes (highlighted in Fig. 3) for which the dual-band backhaul link needs to switch from UB to LB.

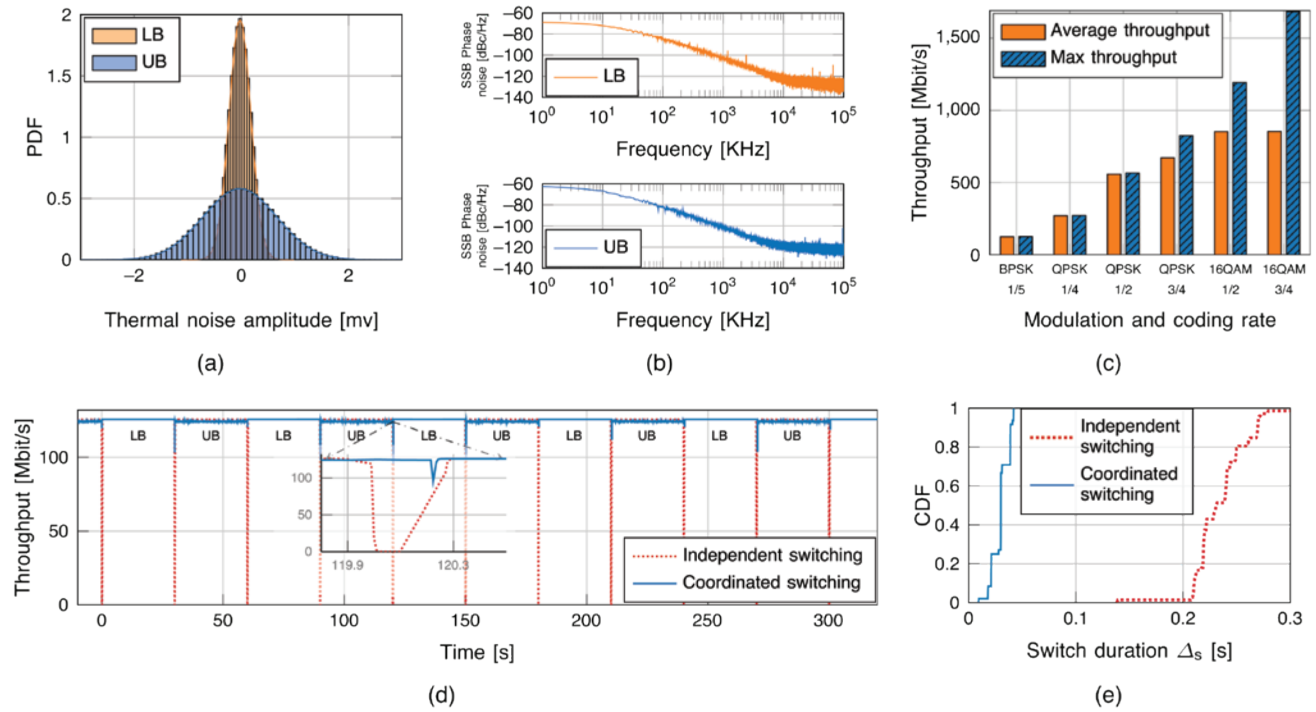

This considered, sharing solutions that would enable contiguous, large bandwidth transmissions in the above-100GHz bands must not be static. Dynamic spectrum sharing would allow regulations to evolve from a firm prohibition of any active transmissions in passive bands, to more flexible schemes where active users could transmit on ultra-large contiguous bandwidths for a portion of time. For example, by considering the 209–241GHz band, only 6.5GHz (non contiguous, 231.5–235GHz and 238–241GHz) are ear-marked for fixed or mobile transmission without the need to share spectrum and thus protect passive incumbents. Figure 4 reports the Shannon capacity as a function of the bandwidth B for different spectrum sharing and aggregation configurations. We consider a received power P_{RX} = – 23dBm, a noise figure F = 8.5dB for the receiver, a noise temperature T_0 = 296K. The Signal-to-Noise-Ratio (SNR) in dB is Γ ¼ P_{RX} 10log$_{10}$(kT_0B) – F, and the Shannon capacity is computed as C = $B\log_2(1+\Gamma)$. According to current regulations, we limit the capacity to the value supported by the maximum bandwidth allowed by the FCC allocations (e.g., 3.5GHz for the "No sharing" configuration). We also consider the possibility of aggregating multiple frequency bands for a single transmission, in this case the "No sharing" configuration would exploit 3.5 + 3 = 6.5GHz. Note that carrier aggregation may not be feasible for all wireless communication technologies. Figure 4 clearly shows the gap in capacity between a system where dynamic spectrum sharing is allowed and one which follows the current regulations, with a gain of up to 8x (no aggregation) and 4.5x (aggregation). Additionally, results in Polese et al. show that, to reach a target rate of 1 Terabit per second (Tbit/s), future 6G systems above 100GHz will require more contiguous bandwidth than what is available without spectrum sharing.

Enabling dynamic spectrum sharing requires (i) awareness of passive users impacted by interference; and (ii) self-adaptive active links. In this paper, we show a possible approach toward gathering awareness of operations of passive sensing instruments in specific frequency, time, and geographical locations. This demonstrates that coordination is indeed

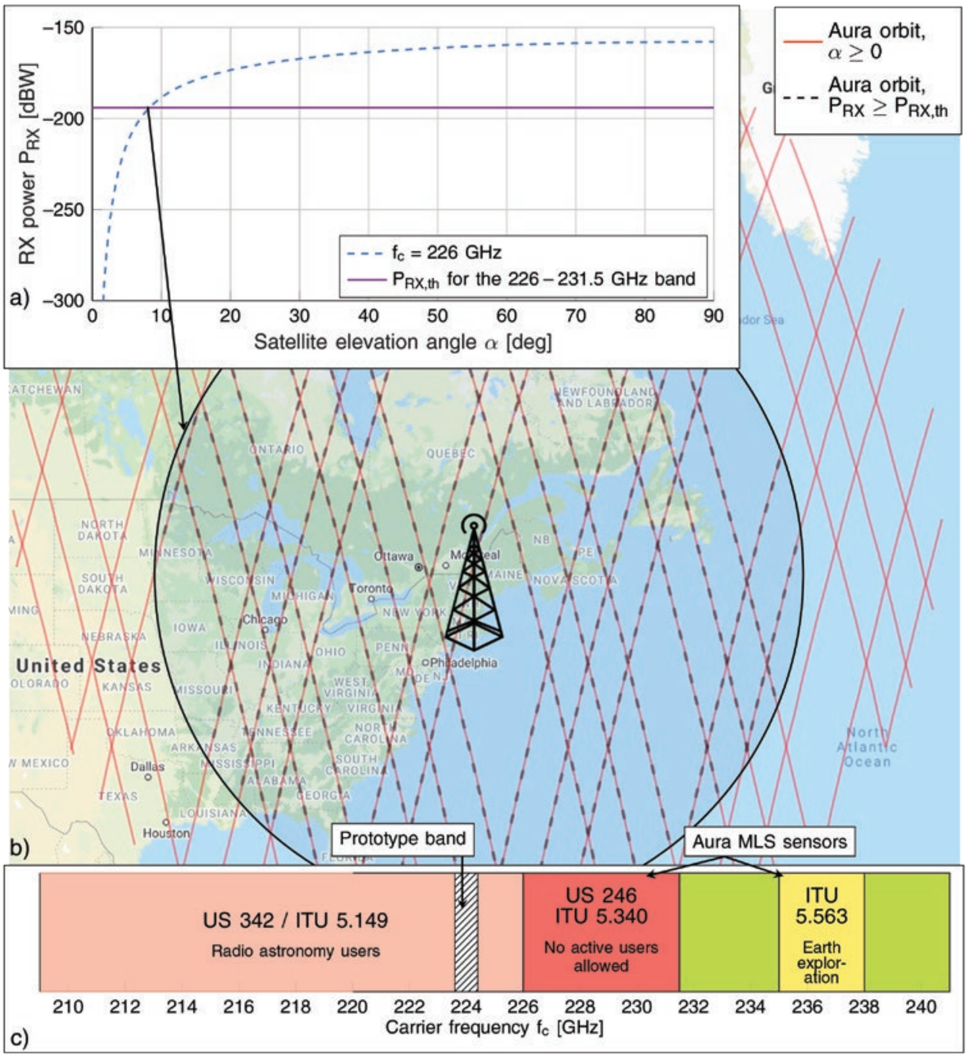

Fig. 2 Experimental result for the dual-band self-adaptive backhaul link. The LB front-end operates in the 123.5–140GHz band (and, in this specific case, at 139GHz), and the UB front-end operates in the 210–225GHz band (specifically, at 224GHz).

a Probability distribution function (PDF) of the thermal noise amplitude for the LB (orange histogram) and UB (blue histogram) front-ends of the dual-band system. Analysis based on *n*=3.2 × 106 noise samples.

b Single Sideband (SSB) phase noise for the LB (orange line) and UB (blue line) front-ends. Analysis based on *n*=1.28 × 10^9 noise samples. dBc is decibel relative to the carrier.

c Average (orange bars) and maximum (blue bars) throughput for different Modulation and Coding Schemes (MCSs). Analysis based on *n* =215,062 throughput samples. Mbit/s is Megabit per second.

d Throughput over time for two independent experiments with MCS Binary Phase Shift Keying (BPSK) 1/5, with independent (dotted red line) and coordinated (solid blue line) band switching.

e Cumulative distribution function (CDF) of duration Δs of a band switch procedure for independent switching (red line) and coordinated switching (blue line). Analysis based on *n* =120 switching events.

feasible, and that scalable systems similar to the Spectrum Access System (SAS) in the Citizens Broadband Radio Service (CBRS) band can be introduced to share information on passive sensing activities and plan interference-free operations with active users.

The results of this paper also illustrate how our prototype is capable of self-configuring in real-time in these frequency bands. Prior literature has shown wireless systems above 100GHz, but only as channel sounders or PHY implementations. This paper evolves the state of the art on above-100GHz communications by introducing dynamic control and dual-band operations. While this prototype focuses on real-time adaptability, its back-end supports a limited bandwidth (i.e., 800MHz), due to constraints in the Analogue to Digital Converter (ADC)/ Digital to Analogue Converter (DAC) rate and signal processing capabilities of its digital fabrics. Therefore, we are also developing a software-defined back-end for multi-GHz-wide transmissions and higher data rates. Nonetheless, the development of a dynamic and real-time re-configurable sub-terahertz link represents a key step to fully unleash the potential of above-100GHz communications.

Next-generation cellular networks will have much higher deployment density, increasing the number of sub-terahertz devices and making manual intervention and configuration close to impossible. A self-organising link improves the resiliency of such deployments. Besides switching operating bands, above-100GHz backhaul links need to automatically adapt to atmospheric conditions, including humidity and wind. Humidity affects the wide-band frequency response, increasing the molecular absorption (and thus the path loss) selectively over the electromagnetic spectrum. Consequently, wireless links need to select the appropriate bandwidth or use multi-carrier systems with adaptive MCS, carrying more information on carriers less affected by absorption. Wind, instead, may impact directional transmissions by introducing vibrations in the radio infrastructure.

Conclusions

Today, the spectrum above 100GHz is rigidly regulated and highly fragmented, with several narrow bands reserved for passive users. This prevents communication systems from exploiting ultra-wide bandwidths.

This article has demonstrated the feasibility of dynamic spectrum sharing in the above-100GHz spectrum. We have introduced a dynamic, dual-band backhaul link, showcased real-time adaptability of an above-100GHz wireless link, and discussed how dynamic spectrum sharing benefits the development of wireless communications systems in these frequency bands.

This fundamental result demonstrates how future regulations for this spectrum band may rely on sharing technologies to guarantee seamless coexistence of passive and active users, and fosters future research on design of passive/active awareness and coexistence schemes in the spectrum above 100GHz.

For more infomration, see https://doi.org/10.1038/s44172-022-00002-x