



WaveFlex: A Smart Surface for Private 5G CBRS Networks

FAN YI, Princeton University, USA

KUN WOO CHO, Princeton University, USA

YAXIONG XIE, University at Buffalo, SUNY, USA

KYLE JAMIESON, Princeton University, USA

We present the design and implementation of **WaveFlex**, the first smart surface that enhances Private 5G networks operating under the shared-license framework in the Citizens Broadband Radio Service frequency band. WaveFlex works in the presence of frequency diversity: multiple nearby base stations operating on different frequencies, as dictated by a Spectrum Access System coordinator. It also handles time dynamism: due to the dynamic sharing rules of the CBRS band, base stations occasionally switch channels, especially when priority users enter the network. Finally, WaveFlex operates independently of the network itself, not requiring access to nor modification of the gNB or UEs, yet it remains compliant with and effective on prevailing cellular protocols. We have designed and fabricated WaveFlex on a custom multi-layer PCB, software defined radio based network monitor, and supporting control software and hardware. Our experimental evaluation benchmarks operational Private 5G and LTE networks running at full line rate. In a realistic indoor office scenario, 5G experimental results demonstrate an 8.58 dB average SNR gain, and an average throughput gain of 10.77 Mbps under a single gNB, and 12.84 Mbps under three gNBs, corresponding to throughput improvements of 18.4% and 19.5%, respectively.

CCS Concepts: • **Networks** → **Mobile networks**; • **Hardware** → **Wireless devices**.

Additional Key Words and Phrases: 5G Network, Metasurface, Smart Surface, Beamforming

ACM Reference Format:

Fan Yi, Kun Woo Cho, Yaxiong Xie, and Kyle Jamieson. 2024. WaveFlex: A Smart Surface for Private 5G CBRS Networks. *Proc. ACM Netw.* 2, CoNEXT4, Article 37 (December 2024), 21 pages. <https://doi.org/10.1145/3696394>

1 Introduction

Starting in 1986 with the original design of the predecessor of WaveLAN, and continuing to 802.11 Wi-Fi, wireless networks deployed by enterprises [30, 39] and home users have operated on a technology trajectory determined by Wi-Fi hardware, lagging years behind its wireless cellular network cousin.

That status quo is poised to change, however, in the context of the “*Industry 4.0*” manufacturing, retail, and general business trend, and specifically with the advent of locally-managed enterprise 5G deployments [40]. These *Private 5G* cellular networks have recently emerged as a serious competitor that may soon supplant traditional Wi-Fi, due to their superior capabilities. Networks that scale in density to millions of nodes per square mile will support massive Internet of Things deployment of embedded devices and tags. Networks that scale latency down to milliseconds will support real-time control applications. And networks that scale data rates and capacity to multi-Gbit/second and Tbit/second aggregate throughput per square mile will fully connect users.

Authors’ Contact Information: Fan Yi, Princeton University, Princeton, New Jersey, USA, fanyi@princeton.edu; Kun Woo Cho, Princeton University, Princeton, New Jersey, USA, kwcho@princeton.edu; Yaxiong Xie, University at Buffalo, SUNY, Buffalo, New York, USA, yaxiongx@buffalo.edu; Kyle Jamieson, Princeton University, Princeton, New Jersey, USA, kylej@princeton.edu.



This work is licensed under a Creative Commons Attribution-ShareAlike International 4.0 License.

© 2024 Copyright held by the owner/author(s).

ACM 2834-5509/2024/12-ART37

<https://doi.org/10.1145/3696394>

Private 5G networks can operate in a number of bands in the radio spectrum, each of which may be accessed under a *licensed*, *unlicensed*, or *shared* basis, according to the relevant government regulations. Shared spectrum access—in the *Citizens Broadband Radio Service (CBRS)* band in the U.S. [17]—is making grassroots, bottom-up deployment of Private 5G networks possible, spurring innovation in a way similar to Wi-Fi in the early 2000s. CBRS offers a unique opportunity for enterprises to deploy Private 5G networks without needing to acquire expensive, exclusive-use licensed spectrum, even further catalyzing innovation.

Private 5G networks in CBRS, while promising, encounter several obstacles, particularly in indoor environments. The relatively high-frequency, high data rate nature of 5G signals makes them susceptible to attenuation and interference. The coverage provided by 5G small cells is often inconsistent and struggles to penetrate through certain materials, leading to dead zones in areas such as basements or building corners [4]. Even within covered areas, there can be locations with low throughput due to physical obstructions or the suboptimal placement of small cells. Existing research, in the context of various other network types (including Wi-Fi and IoT) [7, 15, 26, 54], has demonstrated that incorporating smart surfaces in indoor scenarios can mitigate these problems by dynamically adapting the wireless environment.

This paper describes the design and implementation of *WaveFlex*, the first practical smart surface operating at full network line rate to enhance Private 5G wireless networks in the CBRS band, and sharing that band with incumbent priority users. Referring to Fig. 1, the shared nature of the CBRS band and the complexity of 5G network imposes new design goals:

- (1) **Multi-channel:** First, private 5G networks often comprise multiple base stations (*gNBs*) operating at diverse frequencies from 3.55 GHz to 3.7 GHz and in close proximity, requiring WaveFlex to have the versatility to optimize multiple frequency channels at the same time.
- (2) **Autonomous:** Second, since WaveFlex is not a part of the existing 5G network infrastructure and lacks direct access to both *gNB* and mobile user equipment (*UE*), its control module must function autonomously, without explicit feedback from these entities, yet still compliant with prevailing protocols.
- (3) **Adaptive:** Finally, due to the dynamic sharing nature of the CBRS band, *gNB* operating frequencies may vary, especially when higher-tier users enter the network. WaveFlex must therefore be able to adapt its hardware to these changing frequencies.

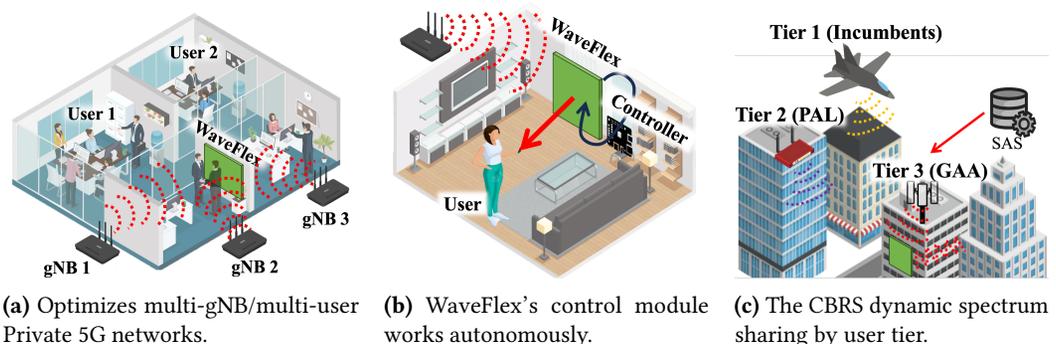


Fig. 1. *WaveFlex*, a surface for CBRS Private 5G/LTE networks with three design goals: (a) Multi-channel, (b) Autonomous, (c) Adaptive.

Through a novel combination of hardware and software co-design *WaveFlex* solves each of the foregoing design challenges. *WaveFlex*'s multi-layer PCB surface design (§4) integrates miniaturized,

custom-designed, and frequency-tunable channel filters with surface-mounted amplifiers to allow the surface to effectively [56] target multiple CBRS channels in a broad frequency spectrum simultaneously. WaveFlex also dynamically adjusts to frequency changes in real time, effectively addressing the spectrum sharing challenges inherent to the CBRS band. Integration of a high-resolution 5G channel monitor with a hardware-software real-time surface controller enables the surface to synchronize with the Private cellular network and adapt its hardware operation at line rate, yet remain independent of (and thus not constrained by) the network itself. Thus, WaveFlex offers deployment advantages, not requiring coordination with or cooperation with both 5G network equipment vendors and network operators, realizing the first steps toward the vision of smart surfaces for the next generation of cellular networks [9].

WaveFlex is backward compatible with Private LTE networks, and so our experimental evaluation (§5) measures the performance gains WaveFlex achieves in each, compared to Private 5G/LTE networks operating without the system. Microbenchmarks characterize an average SNR gain of 8.58 dB at the physical layer enabled by the surface’s amplifier, phase shifter and filter circuitry. End-to-end experiments are conducted to evaluate the performance of WaveFlex in Private 5G/LTE networks operating at line rate. The experiments measure throughput for single and multiple gNB/eNB scenarios, as well as the system’s ability to adapt in real-time to wireless channel and traffic load changes. On average, WaveFlex achieves a 10.77Mbps improvement in downlink throughput under a single gNB. For locations with low throughput (i.e., below 40Mbps), WaveFlex demonstrates an average throughput gain of 12.40 Mbps, resulting in a 1.37-fold increase compared to scenarios without WaveFlex. Multi-gNB experiments demonstrate that WaveFlex is able to achieve 12.84 Mbps throughput gain under three gNBs, corresponding to a throughput improvements of 19.5%.

2 Primer: Private 5G in CBRS

5G TDD PHY Layer. Private 5G networks in the CBRS band typically work in Time Division Duplexing (TDD) mode. TDD operates by segmenting time into slots, with some slots dedicated to downlink data transmissions and others to uplink. For example, Fig. 3 illustrates the uplink-downlink (UL-DL) slot pattern and frame structure for FR1.30-2 [1]. Such a slot pattern exhibits a UL-DL switching periodicity of five slots, each containing 14 OFDM symbols. We also emphasize that there also exists a special slot (S) that helps the transitions between downlink slots (D) and uplink slots (U), which contains a combination of downlink, uplink symbols, as well as guard symbols that provide tolerance to synchronization uncertainty.

The length of each slot depends on subcarrier spacings (SCS). In 5G FR1, there are three available SCS: 15 kHz, 30 kHz, and 60 kHz. These SCS values correspond to slot durations of 1 ms, 0.5 ms, and 0.25 ms, respectively. The choice of numerology depends on the specific requirements of the 5G use case, such as latency, coverage, and capacity. The example in Fig. 3 utilizes a SCS of 30 kHz, corresponding to a slot duration of 0.5 ms.

CBRS Band. The CBRS band for 5G networks is defined as Band n48 by 3GPP, ranging from 3550 to 3700 MHz. Private 5G/LTE networks are unlicensed users, typically operating with bandwidths of 10 or 20 MHz. Real-time coordination and spectrum allocation of CBRS users is managed by a *Spectrum Access System (SAS)*. This complexity, however, introduces dynamism to the network, and further, just like Wi-Fi, successful large Private 5G/LTE network deployments become the victim of their own success—the radio spectrum becomes crowded with multiple base stations (*gNBs/eNBs*) operating concurrently, each operating on distinct *channels* within the CBRS band.

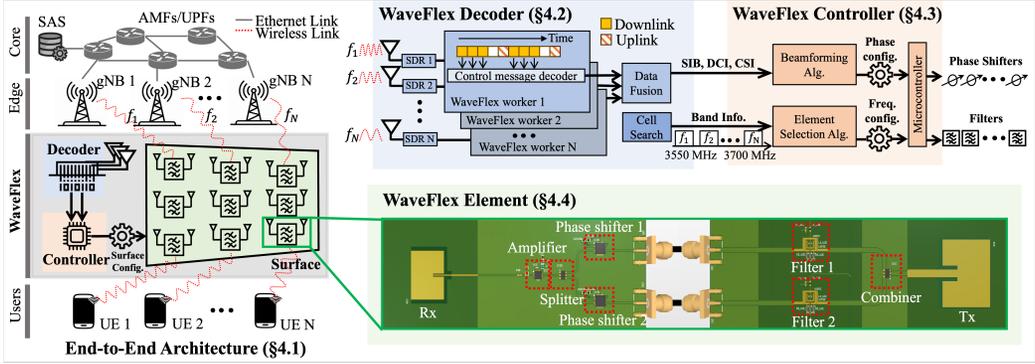


Fig. 2. WaveFlex architecture overview: The **decoder** monitors a private 5G network’s operation and synchronizes the system to that network, the **controller** optimizes the surface’s configuration using data from the decoder, and the **surface** amplifies, filters, and modulates ambient network transmissions.

3 Design

We begin by detailing WaveFlex’s architecture in Section 3.1. Section 3.2 describes WaveFlex decoder, which enables WaveFlex’s autonomy. Section 3.3 outlines the WaveFlex controller in the context of a real private 5G network deployment. Section 3.4 describes our multi-channel hardware, emphasizing the tunable filter design.

3.1 System Overview

Fig. 2 presents the architecture of WaveFlex, encompassing the *decoder*, the *WaveFlex controller*, and our frequency-tunable multi-channel *surface* itself.

WaveFlex decoder. A distinctive feature of WaveFlex is its capability to operate independently without needing explicit information from the SAS, the 5G gNB, and UE. Instead, WaveFlex monitors these components of the 5G network by decoding the control channel messages shared between the gNB and UE. Keeping in step with the time schedule of the 5G network, the WaveFlex decoder synchronizes with the gNB, and then decodes network configurations from System Information Blocks (SIBs), and granular, sub-millisecond-level *Downlink Control Information (DCI)*. The decoded SIBs and DCI are sent to the WaveFlex controller for subsequent processing.

WaveFlex controller. Based on the received SIBs and fine-grained DCI, the WaveFlex controller runs beamforming and element selection algorithms to calculate appropriate phase and filter configurations, respectively, which are subsequently sent to the microcontrollers for real-time adjustments.

WaveFlex surface. The WaveFlex surface (Fig. 2) consists of a series of amplified unit elements, the circuitry of which is illustrated in the lower right of the figure. The surface is designed to operate in both reflective and transmissive modes, providing flexibility in its application. Each unit element operates independently on each of two 5G channels, as described in Section 3.4. An integrated microcontroller applies phase and filter configurations to phase shifters and custom tunable filters, as instructed by the WaveFlex Controller.

3.2 Decoder Design

The WaveFlex decoder performs three main functions. The *cell search* that searches for all nearby private 5G gNBs within the CBRS band; the *configuration decoder* that synchronizes with the detected gNBs and decodes the System Information Blocks (SIBs) to obtain the necessary 5G

network configuration; and the *control channel decoder* that decodes the gNBs' control channel to extract DCI. The list of gNBs, network configurations and extracted DCI are then relayed to the WaveFlex controller for filter control, controller parameter setup and beamforming phase control, respectively.

Cell search. The WaveFlex decoder executes cell search to identify gNBs within CBRS band, which involves traversing the CBRS band and attempting to detect 5G Synchronization Signal Blocks (SSB). The WaveFlex decoder efficiently searches for SSBs using a 5G SA-specific frequency raster¹ defined by 3GPP [2]. If the decoder loses a gNB's signal due to the gNB changing channel, it reinitiates the cell search to reacquire the gNB's frequency.

5G network configuration decoder. 5G offers a high degree of flexibility, with several configurable parameters that impact the operation of WaveFlex controller. Two of the most critical are the *numerology* (which determines the slot duration), and the *slot pattern*. The WaveFlex decoder decodes the System Information Blocks (SIBs) to obtain these values, which are then used to configure the WaveFlex controller, ensuring that WaveFlex works in accordance with the 5G network's specific settings. To stay up-to-date with any configuration changes, the decoder continuously decodes the SIBs, configuring the controller to adapt to variations in real-time.

5G TDD control channel decoder. Prior works [24, 49, 51] have demonstrated the feasibility of an external entity decoding the LTE or 5G control channel to extract DCI. While our 5G decoder is similar in spirit, WaveFlex targets at TDD mode of operation and utilizes DCI for both uplink and downlink data transmissions. The WaveFlex decoder first acquires the SCS and slot pattern from SIB, which defines the arrangement of downlink, uplink, and special slots (Fig. 3). Based on this information, it decodes the control channel for each downlink slot and special slot, skipping uplink slots, since DCI is located only in downlink and special slots. In addition to DCI message decoding, the WaveFlex decoder also computes sub-millisecond level *Channel State Information (CSI)* by utilizing the reference signals transmitted within each slot. The reference signals transmission always exists regardless of whether there is a data transmission in that slot or not.

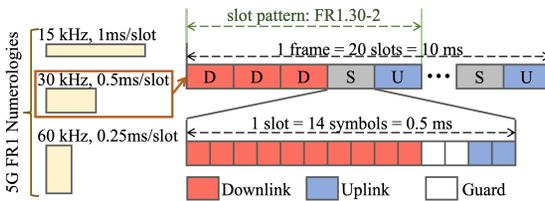


Fig. 3. 5G frame structure, illustrating the TDD duplex mode featuring downlink (D), uplink (U), and special (S) slots in a FR1.30-2 slot pattern frame.

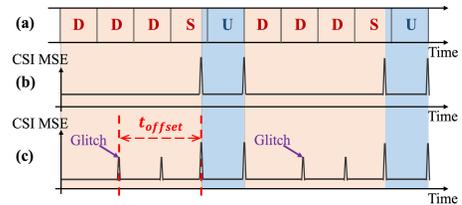


Fig. 4. Operation of WaveFlex's time synchronization: (a) frame structure; (b) synchronized state; (c) unsynchronized state.

3.3 Controller Design

Using the information from WaveFlex decoder, the WaveFlex controller synchronizes with the gNB, calculates optimal phase shifter and filter configurations of each element to maximize data rate.

3.3.1 Synchronization between Controller and gNB. The TDD nature of our target 5G networks imposes challenges on the design of the WaveFlex controller. First, while downlink traffic demand typically exceeds uplink demand by a ratio of 2–3× [52], there is variation between UEs and in

¹In 5G Standalone (SA), the UE does not search for SSBs using the ARFCN raster, as it is narrow and would prolong the search. Instead, 3GPP defines a SSB searching raster with a wider step size for 5G, known as GSCN.

time. Consequently, to support the overall network performance, it is essential to guarantee the reliability of both uplink and downlink traffic. Second, when multiple UEs are active, it is common that the gNB schedules different UEs in different slots, and so WaveFlex must apply different phase shift configurations to optimize different slots, at millisecond time granularity. Referring again to the 5G slot pattern (Fig. 3) the gNB switches from downlink to uplink in the fourth slot of each UL-DL period. The controller needs to synchronize with the gNB and dynamically adjust the phase shifts applied to the surface in accordance with the current link direction.

Controller synchronization. The time-switched nature of TDD places a precise synchronization requirement between the WaveFlex controller and the gNB. The primary objective is to align the switching of phase shifts with the respective gNB's downlink and uplink traffic schedules. To address this synchronization challenge, we employ an iterative offset compensation technique. Initially, the WaveFlex controller's elements switch according to the 5G gNB's periodicity and DL/UL duration. Subsequently, the time offset between the controller's and the gNB's switching time points is measured. Based on this offset, the controller's switching schedule is adjusted to align with the gNB's schedule.

To measure the time offset, our strategy centers on monitoring the periodic pattern of CSI variations. In 5G systems, CSI is collected at the symbol level, as introduced in Section 2. To quantify CSI variations over time, we calculate the mean-squared error (MSE) of CSI values corresponding to adjacent symbols, as follows:

$$\text{MSE} = \frac{1}{N} \sum_{j=1}^N |\text{CSI}_{i,j} - \text{CSI}_{i+1,j}|^2 \quad (1)$$

where $\text{CSI}_{i,j}$ and $\text{CSI}_{i+1,j}$ represent the CSI values at the i^{th} and $(i+1)^{\text{st}}$ symbols, respectively, for the j^{th} subcarrier, and N is the total number of subcarriers.

As shown in Fig. 4, when the WaveFlex controller's element switching pattern is synchronized with the 5G TDD schedule (Fig. 4(a)), we observe a periodic pattern of peaks matched to the TDD schedule as the active transmitter (gNB or any UE) in the schedule changes, as shown in Fig. 4(b). When the WaveFlex controller's phase configuration switching pattern is instead offset by an amount t_{offset} from the 5G TDD schedule, we instead see another periodic pattern of peaks occurring in the CSI MSE curve, which does not match the (known) TDD schedule—labeled *glitches* in Fig. 4(c). These glitches correspond to the CSI changes caused by the downlink/uplink phase shift transitions of the WaveFlex surface. The WaveFlex controller estimates t_{offset} by subtracting the timestamp of the periodic pattern of peaks with the timestamp of the glitch, as shown in Fig. 4(c). The controller then adjusts its switching schedule by $-t_{\text{offset}}$ to converge towards alignment with the network's TDD schedule.

Controller parameter setup. The WaveFlex controller needs to determine the gNB's switching periodicity and DL/UL duration, before it starts synchronizing with the gNB. This information is derived from the slot pattern and slot duration value obtained from the WaveFlex 5G network configuration decoder as mentioned in Section 3.2. The switching periodicity is calculated by multiplying the number of slots in one period by the slot duration. Similarly, DL/UL duration can be obtained by multiplying the number of DL/UL slots in one period by the slot duration.

3.3.2 Blind Beamforming Algorithm. To introduce the WaveFlex beamforming algorithm, we start with the description of the algorithm under ideal conditions with only a single UE and a single gNB. We then extend our algorithm to cover multiple UEs and multiple gNBs.

Our algorithm uses *blind-beamforming*, to run without measuring wireless channels between surface elements to the gNB or UE, [22, 47, 48]. Suppose we have K elements in the smart surface,

and the phases of the K phase shifters are

$$\theta = [\theta_1, \theta_2, \dots, \theta_K]^T. \quad (2)$$

The blind beamforming algorithm searches for the optimal θ^* that maximizes the the channel condition:

$$\theta^* = \underset{\theta}{\operatorname{argmax}} M(\theta) \quad (3)$$

where M (defined below) characterizes the quality of the channel between the transceivers. Blind beamforming solves the problem via an iterative search: it applies a random perturbation $\delta[n]$ at the n^{th} iteration:

$$\theta[n] = \theta[n-1] + \delta[n], \quad (4)$$

where $\theta[n]$ are the phase settings in the n^{th} iteration, repeating the above process until convergence.

Characterizing cellular channel quality. WaveFlex's controller works independently without direct feedback from gNB or UE. Consequently, it cannot directly query the UE for standard channel quality metrics such as the reference received signal power (RSRP), signal-to-noise ratio (SNR), or CSI. However, we observe that the gNBs implement a bitrate adaption algorithm that adjusts the modulation and coding rate to cope with the quality of the channel between the UE and the base stations, making this rate an excellent proxy for assessing cellular channel conditions. More importantly, the gNB broadcasts the rate index and it is thus decodable by the WaveFlex decoder. The WaveFlex decoder serves as a sniffer to decode the bit rate the gNB broadcasts, and calculates the physical data rate R_w , an estimate of the greatest number of bits that can be transmitted over one physical resource block (PRB) without causing errors.

Given the inherent variability of the estimated data rate, we adopt an average rate $\overline{R_w}$ spanning a preset number of slots to mitigate the influence of bit rate index fluctuations:

$$\overline{R_w} = \frac{\sum_{i=1}^{N_{sf}} R_w^i}{N_{sf}}, \quad (5)$$

where R_w^i denotes the rate of the i^{th} slot, and N_{sf} represents the number of slots used for the averaging process. We use this averaged rate $\overline{R_w}$ as the optimization target for a single UE scenario:

$$M_{\text{single}}(\theta) = \overline{R_w}(\theta). \quad (6)$$

Combing Eqs. (3) and (6), we search for the optimal θ that maximizes channel quality $M_{\text{single}}(\theta)$.

Extension to multiple UEs. The blind-beamforming algorithm simultaneously optimizes the channel conditions for multiple UEs. To ensure fairness among UEs, we introduce a channel condition metric tailored for the multi-UE context, which is expressed as:

$$M_{\text{multi}} = \frac{\sum_{j=1}^{N_{UE}} \overline{R_w}(j) \times N_{PRB}(j)}{\sum_{j=1}^{N_{UE}} N_{PRB}(j)}, \quad (7)$$

where $\overline{R_w}(j)$ represents the averaged rate of the j^{th} UE, and $N_{PRB}(j)$ corresponds to the summed number of PRBs allocated to the j^{th} UE in N_{sf} consecutive slots.

Mobility. The mobility of UE imposes stricter requirements on the convergence speed of beamforming algorithms. To address this challenge, we incorporate techniques introduced in BioRARSA [47]. The key feature of BioRARSA is its adaptive nature, which dynamically adjusts the phase shift stepsize based on real-time channel condition feedbacks. This adaptive approach has been demonstrated to improve convergence speed [22, 47, 48].

As illustrated in Eq. (3), the iterative search process maintains a historical record of the highest channel quality, denoted M_{MAX} . This value serves as a benchmark for iteratively searching phase

configurations to optimize the channel condition. In a dynamic environment, M_{MAX} expires when the wireless channel changes, necessitating a mechanism to detect and react to such environment changes. We use consecutive negative feedbacks as a sign of possible channel change. Once the number of consecutive negative feedbacks achieves a pre-defined threshold, N_{nf} , WaveFlex reconfigures phase settings that corresponds to M_{MAX} , and then measure the channel quality again, which obtains M_{new} . If M_{new} is smaller than M_{MAX} , WaveFlex identifies a channel change event, and then resets its beamforming parameters, including M_{MAX} , and $\theta[n]$. Our experimental results in Section 5.2.4 demonstrate that WaveFlex is able to cope with mobility cases.

Convergence Analysis. The convergence of our blind beamforming algorithm can be demonstrated using principles introduced in [47], which reformulates the BioRARSA algorithm as a random search process and then applies the framework proposed by [11] to prove the convergence of BioRARSA.

Our algorithm can be viewed as a series of random perturbations to the phase vector θ , where each perturbation is accepted if it improves the objective function. According to [47] this process satisfies two conditions: firstly, the origin is an interior point of the support of the probability measure for random perturbations; secondly, the objective function $f(\theta)$ is continuous, and all its local maxima are global maxima. These conditions are sufficient to guarantee convergence in probability [11]. Specifically, we can state:

$$\lim_{n \rightarrow \infty} \Pr[\theta[n] \in R_\varepsilon] = 1, \quad (8)$$

where R_ε is the ε -convergence region [28], defined as $R_\varepsilon := \{\theta \in \Theta : f(\theta) > f(\theta^*) - \varepsilon\}$. Here, θ^* is a global maximum point, ε is an arbitrarily small positive number, and $f(\theta)$ is the objective function, $f(\theta) = M_{\text{multi}}(\theta)$. Eq. (8) implies that our algorithm will, with probability approaching 1, converge to a region arbitrarily close to the global optimum as the number of iterations increases.

According to [28, 47], the BioRARSA beamforming algorithm exhibits linear scalability with respect to the number of elements of WaveFlex K . The linear scalability can be expressed as:

$$E[M_{\text{multi}}(\theta[n])] \geq M_{\text{multi}}(\theta^*) - \varepsilon, \quad \forall n \geq c \cdot K, \quad (9)$$

where c is a constant independent of K . This equation implies that the expected number of iterations required to reach within ε of the optimal value scales linearly with the number of elements K . We note that while the number of UEs affects our objective function $f(\theta)$ and may slightly increase the computation time of each iteration, it does not fundamentally alter the computational complexity of the algorithm. We present a quantitative analysis of the convergence speed, and a discussion of the relationship between channel coherence time and beamforming convergence time in Section 5.4.2.

3.3.3 Multiple-gNB element selection algorithm. The two-way element design of WaveFlex surface naturally fits in with two gNB scenarios. For scenarios involving more gNBs, we use a greedy search based element selection algorithm to assign elements to the gNBs. Initially, to ensure fairness among gNBs, we evenly and randomly assign the K elements to the N_{gNB} gNBs and start running the beamforming algorithm. After the beamforming converges, we record the resulting channel quality metric, denoted as M . We then perturb the element selection by reassigning a small subset of elements to different gNBs while maintaining the total number of elements assigned to each gNB. Following another beamforming convergence, the updated channel metric M is compared against its predecessor. Favorable allocations are retained, while detrimental ones are discarded. This iterative procedure is maintained until a stable solution is achieved or until there's a shift in the gNB operating frequencies. To accommodate the dynamism in gNB frequencies, the greedy search is reinitialized based on the updated gNB count N_{gNB} .

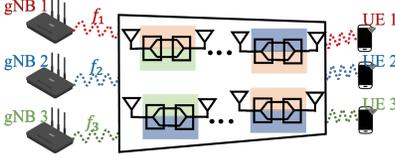


Fig. 5. Multi-channel operation: different elements affect different CBRS channels.

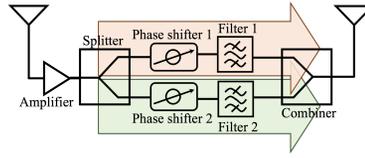


Fig. 6. Circuit schematic of the WaveFlex surface unit element.

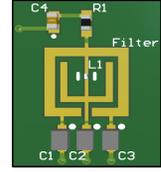


Fig. 7. Tunable filter, the 3D model.

3.4 Hardware Design

Hardware design goal. CBRS networks’ dynamic nature shapes the goals of WaveFlex’s hardware design, which we illustrate in Fig. 5:

- (1) **Frequency tunability:** Private CBRS networks generally involve multiple gNBs working on different *channels*, *i.e.* over non-overlapping frequency ranges. To handle dynamic operating channels, WaveFlex’s hardware must be flexible to adjust its operating frequency.
- (2) **Concurrent cross-channel isolation:** Further, in a multi-gNB network with concurrent, unsynchronized TDD operation at each gNB, WaveFlex’s hardware must support each channel, while not hindering others.

Design of unit element. The WaveFlex *unit element* consists of two *patch antennas* that receive and transmit signal respectively, one *phase shifter* to adjust the transmitted signal phase, an *amplifier* to boost the output signal level, and one *bandpass filter* to achieve our goals of tunability and isolation, as illustrated in Fig. 6. To facilitate multi-channel operation, we integrate two single-channel elements into a unified multi-channel element (Fig. 6). A two-way power splitter divides the received signal and directs each copy into a respective path. The signals produced by the two filters are subsequently recombined using another combiner. Such a configuration enables each single-channel element to independently modulate the signal phase and tune the two respective operating frequencies of each bandpass filter. Our microstrip filter minimizes cross-talk between the two signal paths residing on the same board.

Design of filter. Our filter design goals are as follows: (1) We require a *bandpass* filter that, for each of the two signal paths traversing the unit element, allows the frequency range of a single 5G channel that the controller designates (the *pass-band*) to pass through that path with minimal loss, while significantly attenuating frequencies outside of this range. This allows the path to help the channel of interest without negatively impacting others. (2) Our filter must permit agile adjustment of the pass-band to accommodate the dynamic nature of the CBRS spectrum sharing (see Fig. 1 (c) on p. 2). (3) We require a filter that is cost-effective so that the WaveFlex surface can scale up to large unit element counts, for a realistic deployment.

Resonators are widely employed in filter design due to their frequency selectivity. However, a single resonator often fails to provide satisfactory bandpass performance in terms of insertion loss, roll-off sharpness, and stop-band attenuation. To achieve better filter characteristics, we propose to combine multiple coupled resonators. Specifically, we utilize a dual-mode microstrip filter design [20, 21], which couples two resonant modes with a single resonator structure, offering advantages like space conservation, low insertion loss, and cost-effectiveness. By carefully adjusting the dimensions of the microstrip filter’s metal pattern (Fig. 7), we fine-tune its performance to operate optimally within CBRS band (3.55-3.7 GHz).

To meet our requirements of arbitrarily tuning our operating frequencies, we augment the dual-mode microstrip filter with varactors [21, 46], denoted as $C1$, $C2$, $C3$ in Fig. 7. The varactor is a voltage-dependent capacitor, by applying voltage to it, we are able to change its capacitance, which

affects the filter's operating frequency.² We detail the design of the tunable filter and the derivation of its dimension parameters in Appendix A.

4 Implementation

Surface. Our WaveFlex surface integrates components on custom multi-layer PCBs, as shown in Fig. 8. The surface has 48 elements in total, with an area of $119 \times 76 \text{ cm}^2$. We use SCN-2-35+ splitter [36] and TSS-53LNB3+ bypass amplifier [35] from Mini-Circuits in our unit elements (Fig. 8). We employ Macom MAPS-010144 four-bit phase shifters [33] and an Arduino MKR Wi-Fi 1010 [6] to control the phase of WaveFlex elements at a granularity of $\pi/8$. Our tunable filter is depicted in Fig. 7. For tunable filter control, we use 40-channel AD5370 DACs from Analog Devices[5] to apply variable bias voltages ranging from 2 to 6 V to Macom MAVR-011005-12790T varactors [34].

Decoder. The WaveFlex 5G decoder is modified from NG-Scope5G [49], and its LTE decoder is adapted from NG-Scope [51] to ensure its compatibility with LTE TDD. We pair a laptop with three USRP B210s [16] to execute the WaveFlex decoder (Fig. 9), where each USRP is decoding a distinct gNB/eNB to extract DCI. The extracted DCI provides inputs to our control program, coded in C. This program produces the requisite phase and filter configurations, which are then dispatched to the Arduino microcontrollers and DACs, to adjust the elements' phase shifter and filter.

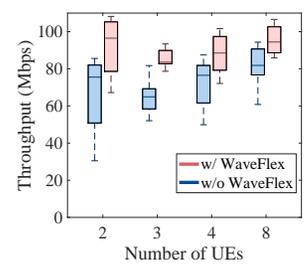
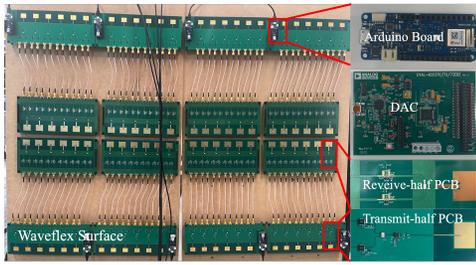


Fig. 8. Hardware implementation: WaveFlex surface combines microcontroller and electronics hardware with custom multi-layer PCB design into an integrated system.

Fig. 9. WaveFlex decoder (upper), private 5G (lower left) and LTE (lower right) small cell.

Fig. 10. End-to-end performance, throughput distribution of WaveFlex under single gNB with 2, 3, 4, 8 UE.

Experimental testbed. We build an operational Private 5G network using three MosoLabs Canopy gNBs [37] and Aether OnRamp 5G core [3, 41]. The three gNBs are configured to operate at different frequencies within the CBRS band. Each gNB has a bandwidth of 20 MHz and uses the slot pattern FR1.30-2, with a slot duration of 0.5 ms. For the Private LTE network, we employ Sercomm CBRS Small Cells [45] and Aether OnRamp 4G core [3]. Motorola Moto G 5G [38] phones serve as the UEs for both 5G and LTE networks. Both networks utilize the Google Spectrum Access System (SAS) [18] to configure their operating frequencies.

5 Evaluation

In this section, we evaluate the performance of our WaveFlex implementation. First, we introduce our evaluation methodology (§5.1). We then demonstrate WaveFlex's end-to-end performance in operational Private 5G network deployment, including evaluations with dynamic mobility and traffic demand (§5.2). We also evaluate WaveFlex's performance under a Private LTE network, covering multiple UE multiple eNB scenarios (§5.3). Finally, we present microbenchmarks evaluating the system under diverse conditions (§5.4).

²This relationship stems from how electrical signals interact with capacitance: capacitance determines how quickly the filter reacts to incoming signals, and this time correlates directly with the filter's resonant frequency.

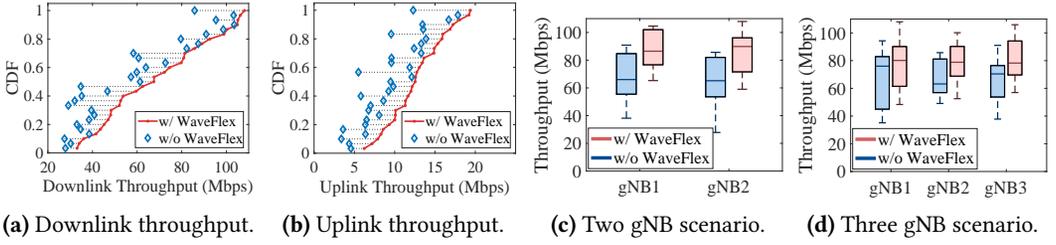


Fig. 11. End-to-end performance: (a) downlink throughput distribution and (b) uplink throughput distribution across 30 locations; and per-gNB throughput in (c) the two gNB and (d) the three gNB scenario.

5.1 Experimental Methodology

End-to-end experiments. We conduct end-to-end communication experiments on both the 5G and LTE testbed (Section 4). Our 5G end-to-end experiments involve three gNBs and eight UEs in total. For the default configuration, the gNBs apply PHY layer settings as described in Section 4. Our LTE end-to-end experiments also involve three eNBs and eight UEs. We use the end-to-end throughput as the primary metric for these experiments, where the throughput is calculated as the total data bits divided by the duration, measured at 100-millisecond intervals.

Microbenchmark experiments. We run experiments evaluating WaveFlex’s robustness across spatial and spectral parameters, and analyzing the control algorithm convergence speed. For our robustness microbenchmarks, we utilize the 5G testbed described in Section 4. Our evaluation metric for WaveFlex’s physical layer performance is the *Signal-to-noise ratio (SNR)* measured at the UE side. To establish a baseline, we measure the SNR in the absence of WaveFlex. Additionally, we take a second baseline by powering on our surface amplifiers, but without executing the WaveFlex controller to optimize performance. We evaluate the system’s performance by comparing SNR when executing WaveFlex, to the SNR collected from these two baselines.

5.2 End-to-end Performance under 5G

In this section, we measure the end-to-end throughput on the 5G testbed under single gNB and multiple gNB scenarios. We also demonstrate the ability of WaveFlex to adapt to real-time mobility and traffic demand changes.

5.2.1 Single UE throughput. We fix the locations of the WaveFlex surface and gNB, and measure both downlink and uplink throughput at 30 different locations to evaluate WaveFlex’s performance in end-to-end private 5G network. As shown in Fig. 16, we perform the experiment in an office scenario of area $12 \times 7 \text{ m}^2$. The gNB (indicated by the triangle in Fig. 16) is located in a meeting room. We vary the locations of the UE (indicated by circles in Fig. 16) to cover 18 locations in the office, and 12 locations in the hallway to include varying distance cases and wall penetration cases.

We show the CDF distribution of downlink and uplink throughput across 30 locations with and without WaveFlex in Fig. 11(a) and Fig. 11(b). On average, WaveFlex achieves a 10.77Mbps improvement in downlink throughput. For locations with low throughput (i.e., below 40Mbps), WaveFlex demonstrates an average throughput gain of 12.40 Mbps, resulting in a 1.37-fold increase compared to scenarios without WaveFlex. These results highlight the effectiveness of WaveFlex in enhancing network performance, particularly in areas with suboptimal throughput.

For uplink, WaveFlex improves throughput by 3.23 Mbps on average. We attribute the lower throughput improvement in the uplink to two main factors. First, in our 5G TDD testbed, the slot pattern is configured as FR1.30-2, allocating $4\times$ time resources to downlink than uplink. This asymmetric resource allocation allows the WaveFlex to have more opportunities to enhance the

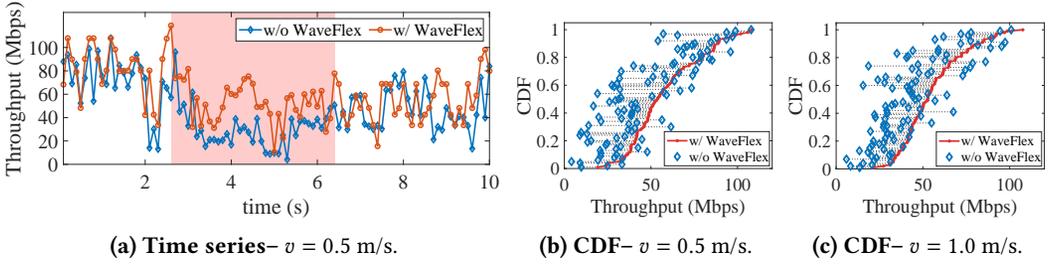


Fig. 12. WaveFlex’s performance under mobility: (a) Throughput trace over 10 seconds with and without WaveFlex at speed $v = 0.5$ m/s; (b)-(c) Throughput distribution with and without WaveFlex at speed $v = 0.5$ m/s and $v = 1$ m/s, respectively.

downlink performance. Second, the gNB transmits at higher power levels than the UE. The higher transmission power in the downlink direction allows the WaveFlex surface to more effectively manipulate and enhance the signal, resulting in better downlink performance improvement.

5.2.2 Multiple UE throughput. We change the number of UEs connected to the gNB to be two, three, four, and eight, for each scenario, we run the experiment 15 times and measure the total downlink throughput of the gNB with and without WaveFlex, to demonstrate WaveFlex’s efficacy. Fig. 10 gives the minimum, 25th percentile, median, 75th percentile, and maximum of the throughput distribution. In a two-UE scenario, WaveFlex improves the throughput by 25.57 Mbps on average. The throughput gain is 15.75 Mbps and 13.38 Mbps when there are four UEs and eight UEs under the same gNB, demonstrating that WaveFlex is able to optimize throughput for multiple UEs simultaneously.

5.2.3 Multi-gNB throughput. In this experiment, we investigate WaveFlex’s capability of simultaneously optimizing for multiple gNBs. We change the number of gNBs to two and three, each gNB is connected by one UE, we run each experiment 15 times and present the per-gNB throughput distribution in Fig. 11(c) and Fig. 11(d). WaveFlex achieves a high gNB throughput gain of 19.13 Mbps and 18.63 Mbps, respectively, with two gNBs, since the two-path hardware design naturally fits in with the two-gNB scenario. When increasing the number of gNBs to three, WaveFlex’s average per-gNB throughput gain drops to 12.84 Mbps. We attribute such drops to the less number of elements allocated to each gNB, when the number of gNB increases.

5.2.4 Mobility. To validate WaveFlex’s performance in mobile scenarios, we show a trace of time-sequenced throughput when the UE moves along a specific route in Fig. 12(a). We fix the gNB and surface locations, and initially the distance between the UE and the surface is 1 m. We then move the UE at a constant speed of 0.5 m/s for 10 s. From around 2.5–6.5 s (highlighted in Fig. 12(a)), the UE to gNB line-of-sight (LoS) path is blocked. We move the UE along this route twice to obtain throughput with and without WaveFlex.

During the whole 10 s process, WaveFlex achieves an average throughput gain of 11.78 Mbps, which shows that WaveFlex is able to keep up with environment changes. From 2.5–6.5 s, WaveFlex improves the throughput by a higher value 13.28 Mbps, demonstrating WaveFlex’s capability to provide higher throughput gain when the LoS path is blocked.

We move the UE along the same route with a speed of 1 m/s for 5 s, and present the distribution of throughput with and without WaveFlex with a UE velocity of 0.5 m/s and 1 m/s in Fig. 12(b) and Fig. 12(c), respectively. The average throughput gains achieved by WaveFlex for the two speeds are 11.78 Mbps, and 10.5 Mbps, respectively. As the UE’s speed increases, the WaveFlex controller

detects channel changes more swiftly, as depicted in Section 3.3.2. This results in fewer iterations for beamforming, subsequently leading to a reduced throughput gain.

5.2.5 Traffic demand changes. In this experiment, we use *iperf3* to control the traffic demand of UE to test WaveFlex’s performance under traffic demand changes. We start by limiting the UE’s traffic demand to 25 Mbps for 5 s, and then increase the traffic demand to 80 Mbps for another 5 s.

We show the time-sequenced throughput when UE’s traffic demand changes with and without WaveFlex in Fig. 13. For the first 5 s, the throughput gain of WaveFlex is 3.55 Mbps, we argue that in this period, the wireless link capacity is enough to support the UE’s traffic even without WaveFlex, because the traffic demand is low. From 5 s to 10 s, WaveFlex improves the throughput by 14.93 Mbps, which proves that WaveFlex is able to identify and react to traffic demand changes.

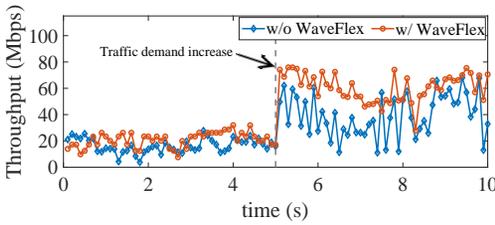


Fig. 13. WaveFlex’s performance under traffic demand change: the throughput trace over a 10-second period, the traffic demand changes at the 5th second.

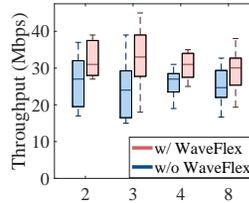


Fig. 14. LTE performance: throughput distribution of multi-UE.

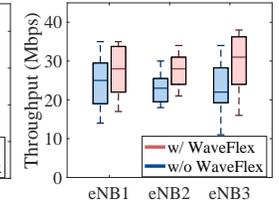


Fig. 15. LTE performance: throughput distribution of 3 eNB.

5.3 End-to-end Performance under LTE

Here we evaluate WaveFlex’s performance on the LTE testbed with varying UE counts and under the three-eNB scenario.

Varying UE count. To evaluate WaveFlex’s performance with multiple UEs, we connect two, three, four, and eight UEs to a single eNB and measure the downlink throughput of the eNB with and without WaveFlex. From Fig. 14, we can observe that WaveFlex achieves an overall average throughput improvement of 6.27 Mbps, including cases with two, three, and four UEs. In the case of eight UEs connected to the eNB, WaveFlex is able to improve the throughput by 3.88 Mbps.

Multi-eNB throughput. In this experiment, we investigate WaveFlex’s multi-channel optimization under the LTE network. In our setup, the number of eNBs is three, and each eNB is connected with one UE. The per-eNB throughput distribution is shown in Fig. 15. On average WaveFlex improves the per-eNB throughput by 2.64 Mbps, 4.06 Mbps, and 5.64 Mbps, respectively, for each eNB.

5.4 Microbenchmark Experiments

In this section, we use our 5G testbed to evaluate WaveFlex’s effectiveness in improving SNR under different conditions, to discern the source of improvements.

5.4.1 WaveFlex Robustness Across Key Parameters. We begin by presenting the Microbenchmark results across spatial and spectral parameters, where the gNB operating frequency is 3567.18 MHz.

Performance across locations. We fix the location of the WaveFlex surface, and the gNB as in Fig. 16, and measure the SNR at 30 different locations to evaluate the SNR improvement of WaveFlex. At each location, the experiment is conducted 10 times. We calculate the average SNR over the 10 measurements and plot them in Fig. 16. The figure shows results under three scenarios, Fig. 16(a) disabling amplifiers and the WaveFlex controller, Fig. 16(b) enabling amplifiers but without

executing the WaveFlex controller, and Fig. 16(c) enabling amplifiers and executing the WaveFlex controller. From this figure, we can observe that WaveFlex provides high SNR improvements. Compared with the environment-only baseline, WaveFlex improves SNR by 8.58 dB for the 30 indoor locations on average. For the 12 hallway locations, WaveFlex can achieve a SNR improvement of 9.08 dB. Comparing Fig. 16(a) and Fig. 16(b), at some locations only enabling amplifiers results in a lower SNR compared to the environment-only baseline case where amplifiers are disabled. Fig. 18(a) shows the CDF distribution of received SNR across 30 locations. On average, WaveFlex outperforms the amplifier case by 7.66 dB.

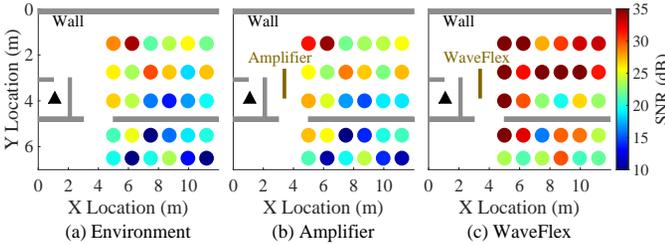


Fig. 16. SNR map at 30 locations under three scenarios: (a) the environment alone, (b) the environment with amplifiers on, (c) the scenario with WaveFlex running.

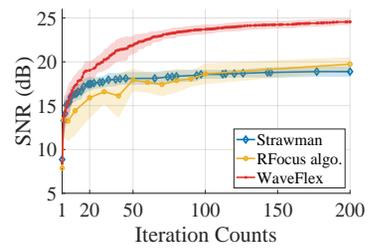


Fig. 17. SNR performance comparison of WaveFlex, RFocus, and the strawman algorithm over iterations.

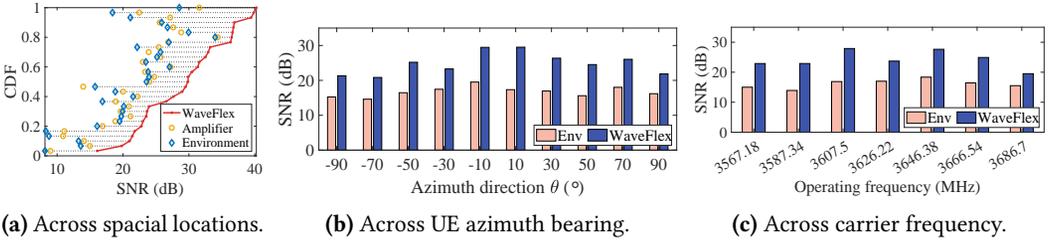


Fig. 18. Microbenchmark SNR values stratified by spacial location, azimuth angle, and carrier frequency.

Performance across azimuthal bearing. We measure WaveFlex’s efficacy under different azimuthal bearings of the UE to the WaveFlex surface. We keep the gNB and the WaveFlex surface location unchanged, and maintain the distance between the surface and the UE to 5 m. We vary UE locations to change the UE to surface azimuthal angle from -90° to 90° with a step size of 20° . Fig. 18(b) presents the average SNR with different azimuth directions with and without the WaveFlex surface. We can see that WaveFlex enhances SNR by 8.10 dB on average across the entire azimuth range. WaveFlex’s lowest improvement in this experiment is 5.70 dB, which happens when at a 90° bearing.

Performance across carrier frequency. To validate WaveFlex’s performance across the CBRS band, we conduct experiments that configure the center frequency of the gNB and the surface to span 3.55 GHz to 3.70 GHz. The center frequency of a 5G gNB cannot be set arbitrarily, it must adhere to a 5G SA-specific frequency raster [2]. We configure the gNB center frequency in steps of approximately 20 MHz, ranging from 3567.18 MHz to 3686.70 MHz, as illustrated in Fig. 18(c). We present the measured SNR with and without WaveFlex, in Fig. 18(c), where WaveFlex enhances the SNR for the entire CBRS band, achieving an average SNR improvement of 8.33 dB.

5.4.2 Control algorithm convergence speed. In this section, we quantitatively analyze WaveFlex’s convergence speed under the 5G testbed environment, and compare it with a baseline algorithm and the RFocus [7] control algorithm. Our baseline algorithm is a strawman blind beamforming approach, which applies random phases to the elements multiple times and selects the set of phase configurations yielding the highest channel quality. In this experiment, we utilize our 48-element surface hardware as a constant factor, fix the locations of the UE, gNB, and WaveFlex surface across all comparisons, and run each algorithm 10 times to ensure a fair evaluation.

Figure 17 illustrates the achieved SNR as a function of the iteration count for the three algorithms under comparison, where the curves represent mean values, while the shaded areas depict the standard deviation across the 10 runs. Our results demonstrate that WaveFlex achieves a high SNR of 24.57 dB after convergence. In contrast, both the strawman algorithm and RFocus fail to achieve the optimal phase configuration within 200 iterations. Additionally, WaveFlex achieves a high SNR gain of 13.48 dB within 50 iterations, which translates to 100 ms, when considering the DCI is collected every 0.5 ms, and the experiment adopts 4 DCI for one iteration. Our experimental observations reveal that the channel coherence time for static UEs typically ranges from 400 to 500 ms, providing ample time for WaveFlex to converge. However, for mobile UEs, this coherence time is significantly reduced, with the exact duration depending on the UE’s velocity. We note that WaveFlex is still able to achieve a high SNR gain of 10.63 dB within 20 iterations. This rapid improvement demonstrates WaveFlex’s robust performance in challenging wireless environments, where the channel could change before WaveFlex converges.

Notably, the RFocus curve exhibits fluctuations and large standard deviations, particularly during the initial iterations. This behavior stems from RFocus’s reliance on element-wise majority voting across random guesses. When the number of random guesses is insufficient, the majority voting mechanism may lead to suboptimal phase selections.

6 Related Work

A smart surface, also known as a *reconfigurable intelligent surface (RIS)*, augments—but does not replace—the wireless infrastructure of base stations and mobile users [8, 27, 31, 50]. Earlier work has described the design of smart surfaces with *passive* elements to improve the performance of Wi-Fi links: LAIA [26], RFocus [7], and ScatterMIMO [15] are representative examples. These systems have demonstrated compelling performance improvements for Wi-Fi, but stop short of integrating real-time, line-rate 5G/LTE operation into their implementation.

Cao *et al.* [10] and Yuan *et al.* [53] design RIS-assisted MAC protocols for Wi-Fi-like networks using a combination of CSMA and TDMA. However, their efforts do not address cellular networks or provide experimental evaluations. Other work explores the use of *active* elements. Long *et al.* [32] describe the architecture, compute link budgets, and propose optimization algorithms for such systems but stop short of hardware realization. LAVA [54] realizes these ideas for IoT networks at 2.4 GHz, adopting a power sensing mechanism to detect active transmissions. Zhang *et al.* [56] show that active surfaces can achieve multiplicative gains over passive surfaces, yet both stop short of system development and integration with a full-stack cellular network.

Zeng *et al.* [55] investigate smart surfaces to harmonize 5G New Radio and Wi-Fi in unlicensed bands, proposing an alternating optimization. LLAMA [12] investigates a smart surface that rotates signal polarization for IoT networks at 2.4 GHz, while VMScatter [29] designs a MIMO backscatter tag and associated algorithms. NR-Surface [23] is a low-power, real-time smart surface for 5G FR2, but requires modifications to the 5G gNB, hindering widespread deployment.

Dual-band smart surfaces [19, 42, 44] have recently gained attention, but existing designs have limitations in satisfying simultaneous multi-gNB/-user optimization, 5G integration, and autonomy (Table 1). Wall-E [14] introduces a dual-band metasurface for satellite networking, but presents only

Table 1. WaveFlex: comparison to existing surface works.

Related Works	Network Type	Multi-frequency	Auto-nomy	H/W Impl.
[7, 15, 26, 29]	Wi-Fi	Single	-	●
LLAMA [12]	IoT	Single	-	●
LAVA [54]	IoT	Single	●	●
NR-Surface [23]	5G FR2	Single	-	●
[19, 42–44]	-	Dual-band	-	◐
RF-Bouncer [25]	Wi-Fi	Dual-band	-	●
CrossFlit [13]	Wi-Fi	Dual-band	-	●
WaveFlex	5G/LTE FR1	Multi-channel	●	●

● = yes; ◐ = partially yes; - = no

a theoretical analysis. RF-Bouncer [25] investigates a smart surface for two ISM bands to expand indoor wireless coverage, but cannot be directly applied in 5G and requires explicit feedback from the receiver. Chen *et al.* [13] introduce a dual-band metasurface for sub-6GHz and mmWave, but it requires signal measurement from the receiver and its operating frequencies differ by more than 20 GHz, posing a less demanding design challenge. Saeidi *et al.* [43] explore frequency-diverse leaky-wave antennas for one-shot detection of multiple wireless nodes, but do not integrate with real-time systems. In contrast, WaveFlex supports fully autonomous surface control, time and frequency varying operations, and line-rate performance, making it well-suited for 5G/LTE infrastructure.

7 Discussion

Potential Enhancements to the beamforming. While our current blind beamforming algorithm based on BioRARSA demonstrates robust performance Fig. 17, there are indeed opportunities for further improvements in terms of convergence speed and accuracy. One promising direction is the integration of machine learning techniques. For instance, a neural network could be trained on historical data to predict optimal initial phase configurations, potentially reducing the number of iterations required for convergence. Another avenue for enhancement is the incorporation of adaptive algorithms that can adjust the search strategy based on the observed channel characteristics. For example, we could implement a meta-algorithm that switches among different search strategies including gradient descent or simulated annealing, depending on the current states of optimization.

8 Conclusion

We have described the design and practical real-world implementation of WaveFlex, the first smart surface that has demonstrated the capability to enhance the operation of real-world Private 5G/LTE networks operating in the CBRS shared spectrum licensing regime. WaveFlex operates autonomously from the core and RAN, easing deployment by obviating the need to interface with cellular providers or even private cellular network system administrators on-site. The autonomous nature of our design opens up new possibilities to incrementally deploy smart surfaces in private 5G/LTE networks on CBRS band, and gives a new direction to the area in general. Also from an engineering perspective, our design breaks new ground for smart surfaces, adapting to shifting traffic demands and the vagaries of the wireless channel in real time.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grant Nos. AST-2232457 and CNS-2223556, and by a Sony Focused Research Award.

References

- [1] 3GPP. TS138.101-4: User Equipment (UE) radio transmission and reception; Part 4: Performance requirements.
- [2] 3GPP. TS138.104: 5G NR Base Station (BS) radio transmission and reception.
- [3] Aether: An ONF Project. [org].
- [4] M. H. Alsharif, R. Nordin. Evolution towards fifth generation (5G) wireless networks: Current trends and challenges in the deployment of millimeter wave, massive MIMO, and small cells. *Telecommunication Systems*, **64**, 617–637, 2017.
- [5] Analog Devices AD5370 40-ch. 16-bit DAC. [pdf].
- [6] Arduino MKR Wi-Fi 1010. [pdf].
- [7] V. Arun, H. Balakrishnan. RFocus: Beamforming using thousands of passive antennas. *USENIX NSDI Symp.*, 1047–1061, 2020.
- [8] E. Basar, M. Di Renzo, J. De Rosny, M. Debbah, M.-S. Alouini, R. Zhang. Wireless communications through reconfigurable intelligent surfaces. *IEEE Access*, **7**, 116,753–116,773, 2019.
- [9] E. Basar, H. V. Poor. Present and future of reconfigurable intelligent surface-empowered communications [perspectives]. *IEEE Signal Processing Magazine*, **38**(6), 146–152, 2021.
- [10] X. Cao, B. Yang, H. Zhang, C. Huang, C. Yuen, Z. Han. Reconfigurable intelligent surface-assisted MAC for wireless networks: protocol design, analysis, and optimization. *IEEE Internet of Things Journal*, 2021.
- [11] C.-C. Chen, C.-S. Tseng, C. Lin. A general proof of convergence for adaptive distributed beamforming schemes. *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 3292–3295. IEEE, 2011.
- [12] L. Chen, W. Hu, K. Jamieson, X. Chen, D. Fang, J. Gummesson. Pushing the physical limits of IoT devices with programmable metasurfaces. *USENIX NSDI Symp.*, 2021.
- [13] L. Chen, B. Yu, J. Ren, J. Gummesson, Y. Zhang. Towards seamless wireless link connection. *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, 137–149, 2023.
- [14] K. W. Cho, Y. Ghasempour, K. Jamieson. Towards dual-band reconfigurable metasurfaces for satellite networking. *Proceedings of the 21st ACM Workshop on Hot Topics in Networks*, 17–23, 2022.
- [15] M. Dunna, C. Zhang, D. Sievenpiper, D. Bharadia. ScatterMIMO: enabling virtual MIMO with smart surfaces. *ACM MobiCom Conf.*, 2020.
- [16] The Ettus Research USRP B210. [url].
- [17] FCC. 3.5 GHz band overview. [fcc.gov], 2023.
- [18] Google Spectrum Access System (SAS). [url].
- [19] J. Han, R. Chen. Dual-band metasurface for broadband asymmetric transmission with high efficiency. *Journal of Applied Physics*, **130**(3), 2021.
- [20] J.-S. Hong, H. Shaman, Y.-H. Chun. Dual-mode microstrip open-loop resonators and filters. *IEEE Trans. on Microwave Theory and Techniques*, **55**(8), 1764–1770, 2007.
- [21] J.-S. G. Hong, M. J. Lancaster. *Microstrip filters for RF/microwave applications*. John Wiley and Sons, 2004.
- [22] S. Jayaprakasam, S. K. A. Rahim, C. Y. Leow. Distributed and collaborative beamforming in wireless sensor networks: Classifications, trends, and research directions. *IEEE Comms. Sur. and Tut.*, **19**(4), 2092–2116, 2017.
- [23] M. Kim, N. Ahn, S. M. Kim. Nr-surface: Nextg-ready μ w-reconfigurable mmwave metasurface. *USENIX NSDI Symp.*, 2024.
- [24] S. Kumar, E. Hamed, D. Katabi, L. Erran Li. LTE radio analytics made easy and accessible. *ACM SIGCOMM CCR*, **44**(4), 211–222, 2014.
- [25] X. Li, C. Feng, X. Wang, Y. Zhang, Y. Xie, X. Chen. Rf-bouncer: A programmable dual-band metasurface for sub-6 wireless networks. *USENIX NSDI Symp.*, 2023.
- [26] Z. Li, Y. Xie, L. Shanguan, R. I. Zelaya, J. Gummesson, W. Hu, K. Jamieson. Towards programming the radio environment with large arrays of inexpensive antennas. *USENIX NSDI Symp.*, 285–300, 2019.
- [27] C. Liaskos, S. Nie, A. Tsioliaridou, A. Pitsillides, S. Ioannidis, I. Akyildiz. A New Wireless Communication Paradigm through Software-Controlled Metasurfaces. *IEEE Comms. Mag.*, **56**(9), 162–169, 2018.
- [28] C. Lin, V. V. Veeravalli, S. P. Meyn. A random search framework for convergence analysis of distributed beamforming with feedback. *IEEE Transactions on Information Theory*, **56**(12), 6133–6141, 2010.
- [29] X. Liu, Z. Chi, W. Wang, Y. Yao, T. Zhu. VMscatter: a versatile MIMO backscatter. *USENIX NSDI Symp.*, 895–910, 2020.
- [30] X. Liu, A. Sheth, M. Kaminsky, K. Papagiannaki, S. Seshan, P. Steenkiste. DIRC: Increasing Indoor Wireless Capacity Using Directional Antennas. *Proc. of ACM SIGCOMM*, 2009.
- [31] Y. Liu, X. Liu, X. Mu, T. Hou, J. Xu, M. Di Renzo, N. Al-Dahir. Reconfigurable intelligent surfaces: Principles and opportunities. *IEEE Comms. Surveys & Tutorials*, **23**(3), 1546–1577, 2021.
- [32] R. Long, Y.-C. Liang, Y. Pei, E. G. Larsson. Active reconfigurable intelligent surface-aided wireless communications. *IEEE Trans. on Wireless Comms.*, **20**(8), 4962–4975, 2021.
- [33] Macom MAPS-010144 four-bit phase shifter. [pdf].
- [34] Macom MAVR-011005-12790T surface mount GaAs tuning varactor. [pdf].

- [35] Mini-Circuits Low Noise Bypass Amplifier TSS-53LNB3+. [pdf].
- [36] Mini-Circuits Power Splitter/Combiner SCN-2-35+. [pdf].
- [37] MosoLabs Canopy 5GID1 gNB. [url].
- [38] Motorola Moto G 5G Cellphone. [url].
- [39] R. Murty, J. Padhye, R. Chandra, A. Wolman, B. Zill. Designing High Performance Enterprise Wi-Fi Networks. *Proc. of USENIX NSDI Symp.*, 2008.
- [40] L. Peterson, O. Sunay. *5G Mobile Networks: A Systems Approach*. [github], 2022. License: CC BY-NC-ND 4.0.
- [41] L. Peterson, O. Sunay, B. Davie. *Private 5G: A Systems Approach*. [github], 2022. License: CC BY-NC-ND 4.0.
- [42] D. Rotshild, A. Abramovich. Ultra-wideband reconfigurable X-band and Ku-band metasurface beam-steerable reflector for satellite communications. *Electronics*, **10**(17), 2165, 2021.
- [43] H. Saeidi, S. Venkatesh, X. Lu, K. Sengupta. 22.1 thz prism: One-shot simultaneous multi-node angular localization using spectrum-to-space mapping with 360-to-400ghz broadband transceiver and dual-port integrated leaky-wave antennas. *2021 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 64, 314–316. IEEE, 2021.
- [44] Y. Saifullah, Q. Chen, G.-M. Yang, A. B. Waqas, F. Xu. Dual-band multi-bit programmable reflective metasurface unit cell: design and experiment. *Optics Express*, **29**(2), 2658–2668, 2021.
- [45] Sercomm Indoor Enterprise CBRS eNB. [url].
- [46] W. Tang, J.-S. Hong, Y.-H. Chun. Compact tunable microstrip bandpass filters with asymmetrical frequency response. *Proc. of the European Microwave Conf.*, 599–602. IEEE, 2008.
- [47] C.-S. Tseng, C.-C. Chen, C. Lin. A bio-inspired robust adaptive random search algorithm for distributed beamforming. *IEEE ICC Conf.*, 1–6, 2011.
- [48] C.-S. Tseng, J. Denis, C. Lin. On the robust design of adaptive distributed beamforming for wireless sensor/relay networks. *IEEE Trans. on Sig. Proc.*, **62**(13), 3429–3441, 2014.
- [49] H. Wan, K. Jamieson. Evolving mobile cloud gaming with 5G standalone network telemetry. *arXiv preprint arXiv:2402.04454*, 2024.
- [50] A. Welkie, L. Shangguan, J. Gummeson, W. Hu, K. Jamieson. Programmable radio environments for smart spaces. *ACM HotNets Workshop*, 2017.
- [51] Y. Xie, K. Jamieson. NG-Scope: Fine-grained telemetry for NextG cellular networks. *Proc. of the ACM on Measurement and Analysis of Computing Systems*, **6**(1), 1–26, 2022.
- [52] B. Yang, W. Guo, Y. Jin, S. Wang. Smartphone data usage: downlink and uplink asymmetry. *Electronics Letters*, **52**(3), 243–245, 2016. doi:<https://doi.org/10.1049/el.2015.3249>.
- [53] J. Yuan, Y.-C. Liang, J. Joung, G. Feng, E. G. Larsson. Intelligent reflecting surface-assisted cognitive radio system. *IEEE Trans. on Comms.*, **69**(1), 675–687, 2021.
- [54] R. I. Zelaya, W. Sussman, J. Gummeson, K. Jamieson, W. Hu. LAVA: fine-grained 3D indoor wireless coverage for small IoT devices. *ACM SIGCOMM Conf.*, 123–136, 2021.
- [55] M. Zeng, X. Ning, W. Wang, Q. Wu, Z. Fei. RIS aided NR-U and Wi-Fi coexistence in single cell and multiple cell networks on unlicensed bands. *IEEE Trans. on Green Comms. and Net.*, **7**(3), 1528–1541, 2023.
- [56] Z. Zhang, L. Dai, X. Chen, C. Liu, F. Yang, R. Schober, H. V. Poor. Active RIS vs. passive RIS: Which will prevail in 6G? *IEEE Trans. on Comms.*, **71**(3), 1707–1725, 2023.

A Tunable Microstrip Filter Design

We design a *microstrip* filter, a low-cost but effective filter that consists of a series of conductive metal traces printed on a PCB dielectric substrate.

The roll-off steepness of such a filter depends on its order.³ A higher-order filter allows a steeper roll-off but increases the design complexity and signal loss. To mitigate this problem, WaveFlex uses a *dual-mode open loop* microstrip filter as shown in Fig. 19. Each of the two constituent resonator loops functions as a doubly-tuned resonant circuit, halving the number of required resonators for a higher-order filter [21] and thereby simplifying the filter design.

Figure 20 illustrates the operating principle of our filter. Consider unfolding the outer loop (Fig. 19) into a straight line (Fig. 20). The loop resonates when its length $l = n \cdot \lambda/2$ where n is an arbitrary integer, and λ is a wavelength. In Figure 20, we visualize the voltage distribution $V(z)$ ⁴ across the unfolded resonator when $n = 1, 2$ (e.g., at its first two resonant frequencies). Since the resonator's two ends have a zero current, the voltage level maximizes at these two locations. Hence, the curvatures of voltage distribution are restricted to two *modes*.

The first mode is called *an odd mode* with n being an odd integer. Here, the center of the outer resonator has a zero voltage. Since the resonator cannot be excited at the location of voltage nulls, the voltage does not pass into the inner loop, and the resonance occurs only in the outer loop. The second mode is *an even mode* with even n and a maximum voltage level at the center. Here, both the inner loading element and the outer loop resonate, giving rise to two separate resonant frequencies. By carefully adjusting the dimensions of the outer and inner loop, we can tune the operating frequencies. We design our resonator such that two resonant frequencies correspond to the edges of the desired pass band.

Filter dimensions. We now present a detailed derivation of tunable filter dimensions. To concurrently and proportionally adjust the resonant frequencies of the odd and even modes by applying a single bias voltage, the tuning rates of both modes must be analyzed. The tuning rate reflects the change in a modal frequency resulting from a variation in the capacitance C_v . Assuming that when C_v changes from C_{v1} to C_{v2} , the odd-mode frequency shifts from f_{01}^o to f_{02}^o , and the even-mode frequency alters from f_{01}^e to f_{02}^e . The proportional tuning rate condition can be mathematically represented as,

$$R_t = \frac{f_{02}^o - f_{01}^o}{C_{v2} - C_{v1}} = \frac{f_{02}^e - f_{01}^e}{C_{v2} - C_{v1}}, \quad (10)$$

where R_t denotes the tuning rate.

The simplified circuit model for the odd and even modes is illustrated in Figure 21. Y_o , Y_e , θ_o , θ_e represents the equivalent admittances and electrical lengths for the transmission line segments in the figure. According to [21], the input admittance for the odd mode is given by,

$$Y_{ino} = j \left(\omega C_v - \frac{Y_o}{\tan \theta_o} \right), \quad (11)$$

and for the even mode, it is,

$$Y_{ine} = j \left(\omega C_v + Y_o \frac{Y_e + Y_o \tan \theta_o}{Y_o - Y_e \tan \theta_o} \right), \quad (12)$$

³The order of a filter depends on the number of reactive components (e.g., a third-order filter requires at least three components: one capacitor and two inductors, two capacitors and one inductor)

⁴The supply voltage is shared within the resonator loop. The sum of the voltages across the loop is equal to the voltage of the supply.

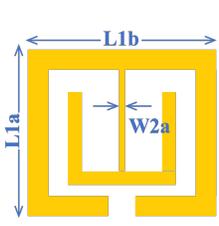


Fig. 19. Filter design.

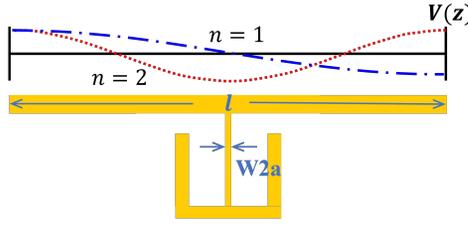


Fig. 20. Voltage distribution in filter.

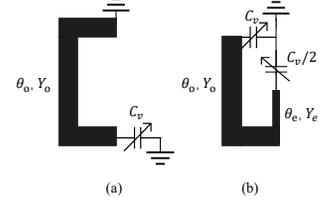


Fig. 21. The model of tunable filter: (a) odd and (b) even mode.

where C_v represents the capacitance of the varactors. The resonant frequencies of the odd and even modes can be determined from the following conditions, respectively,

$$\text{Im}[Y_{ino}] = 0, \text{Im}[Y_{ine}] = 0. \quad (13)$$

By examining Equation 11 and Equation 12, it is apparent that the resonant frequency of the even mode can be altered by adjusting Y_e and θ_e , while keeping the resonant frequency of the odd mode unchanged. We solve the equations 10, 11, 12, 13 to obtain Y_o and θ_o , Y_e and θ_e , that align the tuning rate of the two modes. $Y_{o/e}$ and $\theta_{o/e}$ correspond to the width and length of the metal traces, respectively.

In order to ensure that the adjustable range encompasses the CBRS band from 3550 MHz to 3700 MHz, the resolved Y_e value corresponds to a large inductance value, which equates to a thin loading element that cannot be fabricated. To overcome this issue, we incorporate an inductor $L1 = 0.5$ nH in the loading element, as shown in Figure 7.

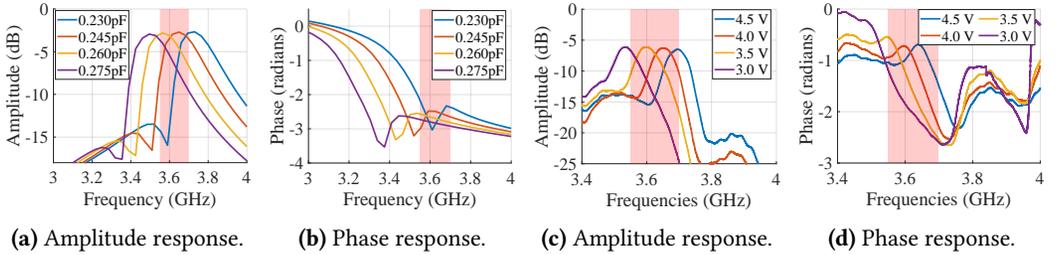


Fig. 22. Simulation and VNA measurement: tunable filter amplitude and phase response versus frequency under varying varactor capacitance and bias control voltages.

B Hardware verification.

This section presents simulated and then empirical (VNA-based) results on our filter performance. In Fig. 22, the pink area highlights the CBRS frequency band, which is our system's target operating frequency. We begin in simulation, Figure 22(a) shows filter roll-off of up to 4 dB per 20 MHz and an insertion loss between -3.5 dB and -3.75 dB when tuning the center frequency within the CBRS band.

To demonstrate WaveFlex's filter response and tunability across the CBRS band, we present our VNA measurements from 3.4 to 4 GHz in Fig. 22(c) and Fig. 22(d). By tuning the bias voltage from 3.0 to 4.5 V, we are able to tune the center frequency from 3.55 to 3.7 GHz. From Fig. 22(c), we see that our fabricated filter exhibits an insertion loss of -6.1 dB, exceeding the simulated insertion loss (*cf.* Fig. 22(a)) by 2.5 dB. We ascribe this increased loss to fabrication imperfections, encompassing factors such as transmission line loss and impedance mismatch. The roll-off steepness is 3.23 dB per 20 MHz.

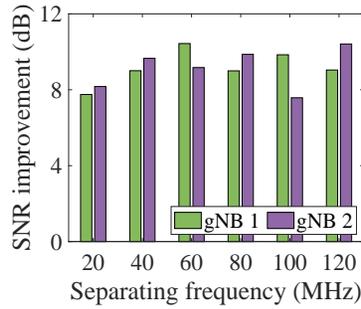


Fig. 23. SNR improvement of two gNBs with different gNB frequency gaps.

C Microbenchmark with different frequency separations

To investigate how the frequency separation of two gNBs influences WaveFlex’s multi-channel performance, we fix the center frequency of gNB 1 at 3567.18 MHz, and change the frequency of gNB 2 from 3587.34 MHz to 3686.7 MHz with a step size of approximately 20 MHz. The SNR improvement of gNBs 1 and 2 is presented in Fig. 23. WaveFlex achieves SNR improvement of 7.75 dB and 8.17 dB for gNB 1 and gNB 2, respectively, when the separating frequency is 20 MHz. At frequency separation exceeding 40 MHz, WaveFlex achieves an average SNR gain of 9.46 dB and 9.33 dB for gNB 1 and gNB 2, respectively.

Received June 2024; revised September 2024; accepted October 2024