Spectraleas: Dynamic Wireless Spectrum Subleasing

Ryan W. West*, M Basit Iqbal Awan*, Miguel Gomez*, Vijay Gopalakrishnan†, Dustin Maas*,
Ashiwan Sivakumar†, Jacobus Van der Merwe*, and Ivan Zelaya†
*University of Utah, †AT&T Labs

***Abstract*— This paper explores the possibility of entities dynamically identifying and subleasing licensed spectrum for short time periods, from other entities that may own spectrum. For example, mobile operators often need more spectrum than available during peak demand, while parts of the available spectrum may go under-utilized at other times when network load is low. We seek to enable dynamic subleasing of the under-utilized spectrum in this situation. To that end, we propose Spectraleas: a system that helps spectrum owners (operators) identify available (needed) spectrum, then matches operators and owners with short-term spectrum needs/availability and initiates co-agreed leases through a spectrum marketplace architecture. While exclusively licensed spectrum will remain the preferred answer for mobile operators, Spectraleas enables spectrum owners to lease spectrum to other entities as needed while maintaining full control over when and where such spectrum is leased. We construct a model of utilized spectrum to take advantage of time varying demands across spectrum owners to allow operators with temporary spectral needs to utilize under-utilized spectrum that other spectrum owners may possess. We identify key challenges that need to be addressed and engineer an architecture to enable such dynamic and temporary sharing of spectrum, down to the level of an individual cell. By analyzing real, active measurements on commercial mobile networks, we find opportunities for spectrum subleasing and show the feasibility of our approach by implementing a prototype mimicking two test operators over a university campus wireless testbed.**

I. Introduction

Mobile Network Operators (MNOs) own spectrum licenses, authorizing them to exclusively operate at certain frequency bands to provide cellular services. However, like most network resources, the spectrum demand curve exhibits spatial and temporal variability. In many instances, an MNO would need additional capacity during peak hours to ease the load off certain bands, while in others, part of the spectrum would go under-utilized [1], [2]. Furthermore, each MNO has different amounts of spectrum and demand by geographical area which adds variability to the utilization of certain spectral bands at different locations and times. In this paper, we propose to utilize this variability in spectrum demand across spectrum owners to enable dynamic subleasing of under-utilized spectrum. This could occur when one MNO who temporarily needs additional spectrum (hereafter *lessee*) subleases from another MNO or spectrum owner (hereafter *lessor*) in a geographic area.

The notion of sharing spectrum is not new. There are known examples of sharing arrangements between MNOs to augment their network capacity to handle demand surges, or swapping portions of spectrum at border locations to increase contiguity [3]. This type of sharing, however, involves arrangements done at long timescales (order of months) through extensive negotiations, planning, and execution. There have also been efforts to share spectrum dynamically among entities using unlicensed spectrum bands. For example, the FCC's ruling on opening the Citizens' Broadband Radio Service (CBRS) band for commercial use enables spectrum sharing between government and other entities via paid non-guaranteed licenses. However, this paradigm requires tight coordination and provides no spectrum availability guarantees. Work on TV Whitespaces and Cognitive Radio (CR) suffers from similar drawbacks, resulting in slow adoption [4], [5].

In contrast to these approaches, we propose Spectraleas, a conceptual system which allows spectrum owners to sublease *exclusively licensed* spectrum temporarily with other MNOs. This dynamic subleasing requires detecting periods of low spectrum use by accurately modeling spectral demand in the past, present, and future. Spectrum owners maintain their spectrum rights, but may allow other operators to temporarily use their spectrum in certain geographical areas in return for some form of compensation. These business angles are out of scope of this paper; Instead, we focus on identifying and addressing three key technical challenges with enabling such short-term spectrum subleasing.

First, a potential spectrum *lessor* must accurately identify available spectrum that can be safely leased, while still providing sufficient coverage and Quality of Service (QoS) for the incumbent's customers. To that end, Spectraleas models cell coverage maps and historical demand to identify pockets of low utilization predicted to persist in the future. Our algorithm allows an MNO to add constraints on the coverage requirements and the minimum QoS guarantees needed before determining leasable spectrum pockets, and considers the risk that potential lessor can use to evaluate the safety of a lease. Our algorithm also covers *lessee* MNO, who should identify areas and duration needing *additional* capacity.

Second, dynamic spectrum leasing requires a framework to communicate and match spectrum supply and demand. We propose a *Spectrum Marketplace*, an entity that coordinates the interactions between lessors and lessees and enables matchmaking via APIs and lease data models, allowing prospective lessors/lessees to indicate interest in or acceptance of inter-MNO spectrum subleases. The Spectrum Marketplace enables MNOs to optimize the utility of their spectrum holdings without changing their license acquisition strategies. Spectraleas tracks active leases on sets of regional base stations (or gNBs) using an end-to-end, automatic deployment process that helps each MNO stay compliant.

Finally, Spectraleas supports mechanisms to reconfigure cellular networks (base stations and cells) for lessors to safely

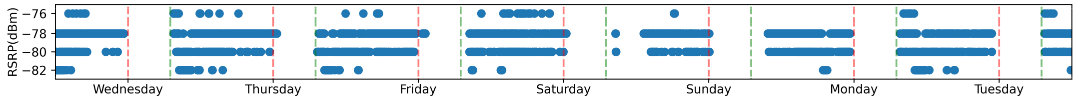

Fig. 1: Sample MNO Cell Uptime Over One Week (red lines indicate 12am, green lines indicate 7am)

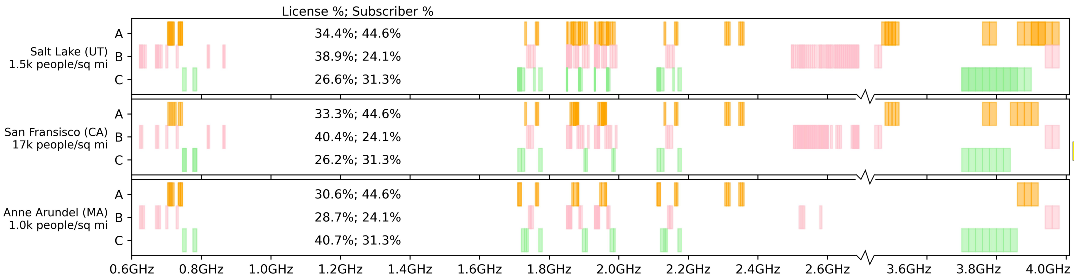

Fig. 2: Spectrum Mid-Band Distribution by County and Licensee (3 MNOs: A, B, and C)

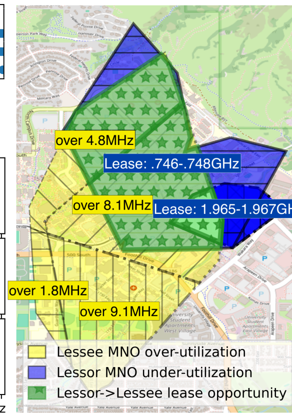

Fig. 3: Supply and Demand

release spectrum which lessees can assimilate to increase overall spectral efficiency. As subleases are intended for short durations, these mechanisms must execute quickly.

We address each of these challenges in this paper. While we target the use case of MNOs leasing spectrum from other MNOs, the proposed mechanisms are generic and can be adapted to other use cases.

We implement a real-world prototype of Spectraleas on a large wireless testbed, the Platform for Open Wireless Data-driven Research (POWDER) [6] using OpenAirInterface (OAI) and Ettus USRP Software Radio devices. We show that our spectrum opportunity algorithm can identify leasable regions between MNOs using historical commercial customer demand data, and that we can release spectrum and deploy it without adversely impacting the performance of the lessor's customers.

Our main contributions include:

1) We present evidence motivating the opportunity for dynamic, pre-owned spectrum leasing among MNOs.
2) We propose an algorithm that geographically identifies under-utilized (or over-subscribed) spectrum pockets.
3) We propose a *Spectrum Marketplace* architecture to manage short-term, dynamic lease transactions. This framework enables MNOs to capitalize (lessor) or supplement (lessee) spectrum holdings, and grants incumbents full control over the timing and location of spectrum leases.
4) We present and evaluate a prototype implementation of dynamic short-term spectrum leasing using 5G Cores and RANs (Radio Access Networks).

II. Motivation

In this section, we use publicly available data and experiments to motivate the possibility of dynamic, short-term spectrum subleasing. We first present evidence of operators putting radios to sleep during periods of low use to save energy and that different operators own varying amounts of spectrum in different geographic areas, indicating that the spectrum could be, in theory, used by other operators that might need it. We preview our evaluation results showing a jointly beneficial instance where a real commercial MNO could have leased spectrum to another commercial MNO. We then show through experiments that operators can provide their users with good performance despite subleasing portions of their spectrum. Finally, we discuss use cases.

Periods of low spectrum use: Traffic demands vary by time, location, season, etc. There is evidence that MNOs already exploit time-varying spectrum demands to power down under-utilized cells during off-peak hours and lower energy bills [7], [1]. Indeed, our own active experiments indicate this phenomenon. We measured the Reference Signal Received Power (RSRP) for one frequency of a base station over one week in late 2023. Figure 1 shows that we detected most coverage in that frequency from 7am until 12am, indicating that the transmission was otherwise powered down due to low demand and presenting an opportunity for MNOs to allow others limited spectrum use without affecting customers.

Diversity in spectrum ownership: Figure 2 shows active cellular spectrum licenses of three large MNOs in United States across three counties, each with varying population densities. This public FCC data [1] shows that MNO A owns licenses for 213MHz of spectrum between 3.4-4GHz in Salt Lake, 163MHz in San Fransisco, and 61MHz in Anne Arundel. In contrast, MNO C owns 201MHz in Salt Lake and 161MHz in Anne Arundel. MNO B has higher holdings between 1.7–2.7GHz in Salt Lake and much lower in Anne Arundel. Figure 2 also shows the percentage of spectrum frequency in MHz that each MNO owns relative to the other MNOs in that county. MNO B owns the most licenses in Salt Lake and MNO C owns the least, but the opposite is true in Anne Arundel. As we cannot access actual MNO traffic demand, we assume each MNO has a uniform share of subscribers in each county as they have nationally. Using the mid-2022 US census and each MNO's publicly reported subscriber base, we annotate Figure 2 with estimated user base as a proxy for demand. Despite the estimations, the data's key takeaway is regional differences in capacity (i.e., amount of spectrum owned) and demand across MNOs, resulting in scenarios where MNO supply and demand may not be evenly distributed.

Measurements of Lessor–Lessee Supply and Demand: As will later be thoroughly explored, we have developed and tested a spectrum opportunity algorithm for Spectraleas that

[1] wireless2.fcc.gov/UlsApp/UlsSearch/searchGeographic.jsp, 12/2023

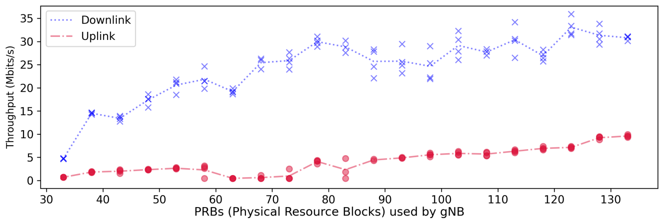

Fig. 4: Throughput & As gNB Bandwidth Decreases

analyzes an MNO cells' historical customer traffic to find geographic regions where that MNO's spectrum may have been under- or over-utilized, and such regions may be good candidates to sell or buy spectrum leases. Our POWDER implementation collected real-time data from several dozen cells from several commercial MNOs. Figure 3 shows a yellow region of cells' coverage area in which one MNO experienced estimated spectrum over-utilization, a blue region where a different MNOs' cells experienced estimated spectrum under-utilization, and a 1.35km^2 green region overlapping the blue and yellow where a temporary lease of 2-4MHz of spectrum could have potentially benefitted both MNOs. The process to arrive at this figure is thoroughly detailed in Sections IV-A and V-A.

Performance impact of subleasing: We measure throughput and latency of a UE connected to a gNB while varying the gNB's total available bandwidth. We run an OAI 5G Core and over-the-air RAN at 3550-3600MHz with 30kHz Subcarrier Spacing (SCS), giving a bandwidth up to 133 Physical Resource Blocks (PRBs). 3GPP specifies few channel bandwidths supported by both UEs and gNBs, so we modify OAI to selectively block an increasing number of PRBs from usage in uplink and downlink, then repeatedly run the gNB with different PRB amounts allowed out of 133 under isolated, indoor channel conditions as Figure 4 shows. A second 'traffic' UE simultaneously sends 5 Mbps of traffic. As PRBs drop from 133 (50MHz bandwidth) to 51 (20 MHz), we observe even latency (not shown) but gradual decline both uplink and downlink throughput. While this is expected, the decline is minimal (partially because fixed control resources are sent regardless of bandwidth), suggesting that reduced bandwidth can still serve existing demand. We ran the same experiment at 3450-3500MHz and in outdoor environments and observed the same trends[2]. Studies such as by Yang et al. [1] have also measured nighttime/weekend periods in which MNOs reduce gNB power for up to 12 hours without impacting user throughput.

Use Cases. The primary use case considered in this work is dynamically subleasing spectrum between two commercial MNOs at certain locations and times when there is disparity in the utilization of different frequency bands across MNOs. However, other use cases could benefit with spectrum subleasing at different time scales and can trivially be accommodated by the system presented in this work.

[2]While absolute throughputs/latencies may change by RAN configuration, the trends are likely to persist and have persisted in every test scenario.

For example, Non-Public Networks (NPN, or private 5G) are often used for industrial use cases such as smart factories/farms and primarily use unlicensed or shared spectrum such as CBRS, but NPN research recognizes the benefits of licensed spectrum [8], [9]. More generally, the spectrum needs for applications involving incremental backups or sensor data reports for remote IoT devices may be predictable, periodic, delay-tolerant, and fine to transmit only at times of low consumer traffic (e.g., early morning subleases each Sunday). Spectrum leasing also further expands the market for entities who own and rent out cell tower space to major MNOs, who could lease spectrum from MNOs to run independent 'utility' 5G networks that connect to MNOs or NPNs using e.g. roaming for supplemental capacity. Finally, a conference or a gaming convention where one MNO having more users could lease under-utilized spectrum from another MNO for a few hours.

III. Related Work

Many studies find fragmentation and varying utilization of 5G bands, suggesting that licensed capacity and demand do not always line up [10], [11], [12]. Some take advantage of these temporarily under-utilized cells by powering down cells [1], [7], while others consider spectrum leasing. The 2023 survey by Parvini et al. overviews spectrum leasing and its economics and explicitly says "the current literature does not provide a clear picture of the implementation perspectives of the various spectrum sharing methods" [13]. It covers CBRS, CR, and Licensed Shared Access (LSA) as major spectrum sharing efforts involving incumbent users with priority who allow secondary users to transmit only during radio silence. Despite much research attention [14], [15] and real-world operations, LSA and CR lack widespread deployment due to the lack of scheduled, guaranteed availability that commercial use cases require [5]. The FCC's CBRS uses a Spectrum Access System (SAS) that allocates two priorities of leases for paying license holders and unpaid general access, and both may be overridden by governmental incumbent use. SAS uses Environmental Sensing Capability (ESC) to detect incumbent signal during leases and may immediately cancel leases, unlike our predictive spectrum identification and guaranteed access of any licensed bands (vs. CBRS's 3550-3700MHz). Neither does CBRS consider how lessee MNOs integrate with the SAS, deploy leases, or consider lease boundaries.

To our knowledge, most studies regarding inter-operator spectrum markets focus on their pricing economics rather than feasible operating logistics [16]. For example, Mohseni et al. develop an economic model between two MNOs that maximizes a lessor's profit while maintaining a sufficiently high QoS [17], but does not identify spectrum opportunities or actually enforce leases. Similarly, Gomez et al. argues that liquid, secondary spectrum markets are necessary to better allocate spectrum vs. governmental spectrum license auctions and propose a model that pools and virtualizes spectrum to improve market interest and usability, which supports our approach but does not work towards deploying it [18], [19].

Wu et al. model a decentralized spectrum sharing system using blockchains to create a Dynamic Spectrum Sharing Architecture and introduce high- and mid-level management servers. These are similar to the Spectrum Liaison and Spectrum Allocator used in our work (See Section IV-C) [20]. However, they focus on decentralization, anonymity, and incentives instead of identifying spectrum opportunities and deployment. Li et al. design a theoretical blockchain allowing multiple lessors/lessees to post intents to buy or sell spectrum and find the optimal buy-sell strategy using a multi-leader multi-follower (MLMF) Stackelberg game model [21]. Like past studies, the contribution here is more in economics than in the technical challenges of spectrum transfer.

Other notable research includes Munoz et al. [22] who devise an architecture involving a market of tenants who can request slice templates similar to ETSI Generic Slice Templates [23]. While this is similar to our lease definitions, they focus on RAN hardware network slicing and network planning. Both papers simulate but do not implement a prototype of their model nor consider integration details. Another system, NeutRAN, uses O-RAN functions in a neutral host framework that automatically deploys multiple tenants' RAN and dynamic spectrum needs onto shared hardware [24]. However, NeutRAN's model does not focus on how spectrum is being acquired or shared between MNOs with separated hardware, and evaluates all deployment requests immediately.

Ultimately, the existing spectrum sharing research does not sufficiently consider the scenario of subleasing pre-licensed spectrum, including identifying potential spectrum opportunities at a technical level and how two MNOs can agree to lease spectrum and seamlessly reconfigure their networks to enforce the lease. Our design seeks to fill these research gaps.

IV. Design

The complete life cycle of Spectraleas's dynamic spectrum leasing consists of A) identifying lease opportunities, B) advertising these opportunities to find interested MNOs, C) negotiating between two interested parties until a lease agreement is reached, and D) applying said lease agreement on all applicable base stations during the lease term.

A. Identifying Lease Opportunities

The key to enabling dynamic spectrum subleasing is identifying spectrum that can be safely leased. We thus spatially model past customer demand from cell coverage area and propose an algorithm that uses coverage maps, historical traffic, and available spectrum to estimate areas where spectrum can be leased without affecting the lessor's QoS or coverage for its customers. If the identified area is large enough for a lessee to transmit without interfering with the lessor, the frequency range and area may be viable candidates for a lease.

1) Spectrum Opportunity Algorithm Overview: Consider 3 types of geographical maps pertaining to a lessor MNO over a specific region of interest (ROI):

1) A map of exclusively licensed spectrum (licenses are often sold at County or Partial Economic Area scale)
2) A service/coverage area map (or all active base stations' coverage area that exceeds an average threshold of at least X RSRP/RSSI, and their operating frequencies)
3) Demand maps of current and historical number of connected UEs per cell, each over avg. 1-15 min. period

These maps can be combined to represent spatial customer demand relative to spectrum, and the algorithm processes them with the ROI to identify potential *partial lease listings* (see Figure 5), which include lease boundaries, frequency ranges, and start/stop times. These listings could be offered on a Spectrum Marketplace (described in Section IV-B) to potential MNO lessees. Each listing is assigned a risk factor, indicating the likelihood that leasing the spectrum would impact the lessor's operations beyond an acceptable threshold, helping to decide whether and how to offer the spectrum.

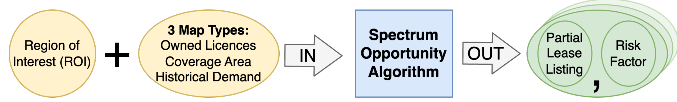

Fig. 5: Spectrum Opportunity Algorithm Inputs & Outputs

The algorithm applies the three maps to the ROI one at a time to identify potential lease listings. This is exemplified in Figure 6 which includes 3 large coverage cells (A and B divided into 120° sectors) and 2 small capacity cells with varying mock customer data. First, it analyzes the preexisting, exclusively licensed spectrum within the ROI, restricting leasable frequencies to those licensed for the region, which may vary if, e.g., the ROI crosses county boundaries. Second, the coverage map is overlaid, and geographic areas outside base station coverage are excluded (e.g., rural or mountainous regions, regions lacking hardware cell coverage). Unused but licensed areas, like those shown in the right edge of zoomed-in Figure 6.2, are less typical but easily leasable barring any regulations, and may be useful for specialized use cases such as experimental research, or smart factories/data centers/farming. Third, the algorithm applies maps of historical UE count and cell utilization to identify areas with consistent high capacity relative to demand (such areas likely cover greater population and are more interesting for lessees). Historical traffic patterns (including recent data to react to ephemeral opportunities) are combined and analyzed for every active cell that overlaps the ROI to find sufficiently large areas with a predictably high capacity relative to demand, as detailed below and in Figure 6.

2) Measurements for Customer Traffic/Capacity: To understand this third step, customer traffic must be measurable, both in terms of bandwidth needed to satisfy a single customer and the traffic that a cell can tolerate based on its bandwidth (or *capacity*). However, one cannot simply observe the proportion of a cell's PRBs occupied in a frame to obtain the cell's remaining capacity, as one cell may schedule the same quantity of PRBs for many mostly idle UEs as another cell for a lone UE downloading and uploading high-throughput VR video, and both are 'capable' of accommodating more traffic.

We thus restrict the term 'capacity' to indicate the maximum number of UEs for which a cell can schedule an average of at least q Hertz of spectrum 'dedicated' to each UE

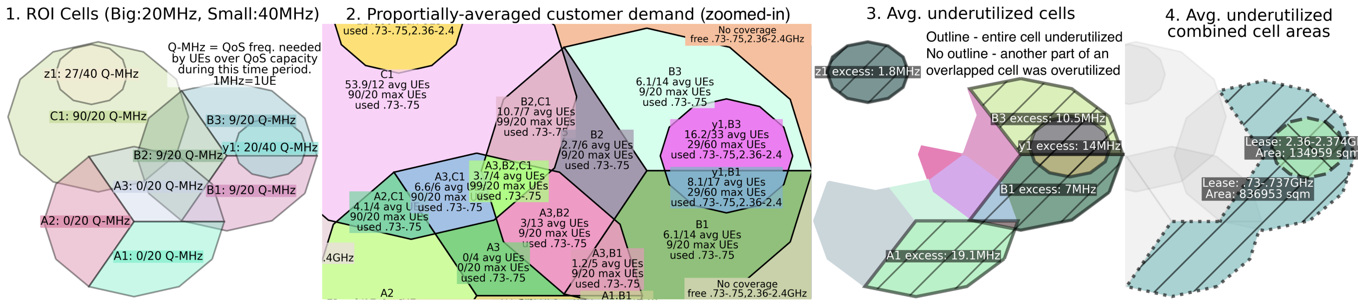

Fig. 6: Spectrum Opportunity Algorithm Process with Example Data

simultaneously, where q satisfies some minimal QoS metric. Guaranteeing a minimum throughput per UE simultaneously is impractical because the number of PRBs required to achieve a certain bitrate varies with channel conditions, Modulation and Coding Scheme (MCS), number of retransmissions, etc. Requiring a certain number of PRBs per user is also not ideal, since a PRB count can represent different portions of overall cell bandwidth, depending on the 5G numerology. For these reasons, we require a minimum average bandwidth (i.e., piece of the spectrum available to a cell) for each active UE, which is independent of UE-specific channel conditions and 5G numerology. This is a conservative minimum as UEs are often idle, and as explored later, risk-tolerant MNOs may reasonably estimate how many UEs out of all connected are actively transmitting to determine cell capacity.

Figure 6 depicts the process of analyzing the customer traffic and capacity of one past time measurement period, and 6.1 shows the number of UEs (or minimum required frequency allocation) for each cell over their total capacities. In this example we assume a UE under *ideal* channel conditions can achieve 1Mbps per MHz and set 1MHz as the minimum QoS target per UE, allowing basic video streaming at peak times (a conservative ratio based on testing). This allows us to calculate a cell's maximum UE capacity that can be simultaneously supported at this QoS target and directly correlate that with the cell bandwidth (e.g., a 20MHz cell can accommodate up to 20 UEs before risking a dip in minimum QoS for any UE during peaks where *all* UEs are simultaneously transmitting). Figure 6.1 shows that cell C was over-capacity while A and B had zero or few connected UEs. This estimated target bandwidth per UE should be changed depending on MNO QoS goals. Thus, we have defined two traffic metrics:

1) Minimum target QoS - the minimum spectrum devoted to an active UE on a cell to reach an acceptable QoE under ideal channel conditions (e.g. 1MHz for 1Mbps).
2) Cell UE QoS capacity - the maximum number of UEs that a cell can accommodate in the peak event that all UEs are simultaneously transmitting at the minimum target QoS.

MNOs may create corresponding Key Performance Indicators (KPIs) to ensure that all active UEs on a cell can hit the minimum target QoS 99.99% of the time. As a cell is highly unlikely to reach its capacity via a peak as defined above due to idle UEs, MNOs may be comfortable exceeding this maximum as long as the probability that a simultaneous UE traffic peak approaches actual Cell UE QoS capacity stays low.

3) Finding Low Spectrum Utilization Pockets in Historical Traffic Maps: Spectraleas's algorithm can use these customer traffic measurements to identifying low-risk spectrum leasing opportunities. It first searches the ROI for areas where historical cell customer demand is low relative to the available spectrum or capacity, known as spectrum pockets. A base station 1+ cells where *cell* here refers to a coverage region transmitting at a single frequency range that cannot be further geographically subdivided (so a typical 3-sector cell here equals 3 cells). However, cell frequency range *can* be contiguously increased/decreased. A cell is likely to experience similar demand regardless of the precise location of its customers, and thus we assume that a past demand measurement period for a cell is on average uniformly distributed throughout its geographic coverage area, and its UEs are compatible across all licensed frequencies (this approach is efficient and improves scalability). An additional, more precise approach for known, non-uniform cell coverage areas is covered later (e.g., a rural freeway has few UEs in adjacent fields).

If a cell's UEs are geographically uniformly distributed, we can divide the ROI's cells along coverage area borders into polygons representing areas with exactly 0, 1, or more overlapping cells as shown in Figure 6.2. Coverage areas overlap when cells of different frequencies overlap (e.g. coverage and capacity cells) or when neighboring same-frequency cells overlap to avoid coverage gaps, creating many unique polygons with varying demands and capacities. To calculate a polygon's demand and capacity for a past time period, we calculate the geographic proportion of each overlapping cell relative to that cell's entire coverage area and separately multiply that against each of the cell's capacity (or frequency bandwidth) and number of actively connected UEs. We then sum all cells' proportional demands and non-overlapping capacities to obtain an overall area-proportional average spectrum utilization and capacity for that polygon (e.g. pentagonal polygon from B2 and C1 in Figure 6.2 could have had a maximum of 90+9=99 UEs (see Figure 6.1) within it if *all* of both cells' UEs had congregated there, but more likely only a geographic proportion of B and C's traffic and QoS capacity (9% and 27%) would be here. B2 thus influences this polygon more and the estimated number of UEs is thus almost within QoS capacity despite C1 being extremely over-capacity).

If this average polygon demand is below its average capacity, the difference represents the polygon's 'unused' capacity during the measurement period. By averaging unused capacities across all of a cell's polygons, we can estimate the how much bandwidth could have been reduced while still satisfying the minimum QoS metric during peak traffic as Figure 6.3 depicts (ranging from 0MHz to the full cell bandwidth). Note that one polygon in a cell can be over-capacity while remaining polygons are under-capacity, causing an under-capacity cell average. Risk-adverse MNOs could reject this cell, but the geographically uniform UE estimation reduces this concern, and a better approach may be to slightly decrease the overall cell capacity to compensate, or potentially ignore it.

Because it is unlikely that all connected UEs were simultaneously transmitting at their minimum QoS bandwidth, an MNO may have a QoS KPI that allows ephemeral periods of 'over-QoS-capacity' (e.g. 0.001% of the time) to allow larger leases at a risk. Figure 6.3 also shows various non-outlined polygons, which represent portions of cells that individually had extra capacity, but are unusable since at least one of their parent cells was over-QoS-capacity elsewhere in the cell.

When a cell has predictable, non-uniform regions of UE traffic, these regions can be divided from the cell and treated like smaller cells (from the algorithm's perspective) as long all intra-cell areas individually have a uniformly-distributed demand. This more accurately represents customer demand and can be calculated with other cells without modification.

When calculating polygon demand, care must also be taken to not double-sum the capacity of same-frequency overlapping cells (common in some MNO deployments), which use inter-cell interference (ICIC) avoidance techniques to simultaneously transmit in one polygon with the 'same' capacity.

4) Combining Low-Utilization Areas into Useful Leases: A lease is only useful if it is large enough for the lessee to transmit within the lease boundaries without interfering with the lessor and vice versa. Therefore, various under-utilized cells with frequency *and* geographic overlap (and optionally 0-cell unused polygons) are combined to form large enough areas to be useful to the lessee without lessor interference, as shown in Figure 6.3 (e.g. a circle with some minimum area (e.g., 200m radius) fits in the combined area). Similarly, a lease must have some minimum bandwidth (e.g. 1.5+ MHz) for meaningful data transfer and duration long enough for gNBs reconfiguration and customer utilization. Figure 6.4 depicts this final step of combining under-utilized cells to form multi-cell areas that could form a lease listing (excluding small cell z1 whose combined smaller size and bandwidth may be less lucrative).

The algorithm must also search across time for candidate pockets of available spectrum that are steady over some minimum duration (e.g. 5 minutes), bandwidth, and area.

5) Additional Search Factors: A trivial process might search for low-traffic points at every e.g. 20 feet of latitude/-longitude in the ROI across all time maps to obtain at least $O(n^3)$ time complexity. By dividing the ROI into overlapping coverage area polygons and assuming uniform cell traffic distribution, complexity here reduces to $O(n)$ as polygons are pre-calculated and only their traffic averages are recomputed at each time period map. Utilization of the Sort-Tile-Recursive algorithm further improves search efficiency [25].

Additionally, a lessor might always provide minimal connectivity in all coverage areas (e.g., LTE's 1.5MHz minimum), even in regions with historically low customer traffic. Only areas with overlapping capacity/coverage cells could fully turned off and lease an entire cell's bandwidth, as at least one remaining cell stays available. Riskier leases may allow a smaller minimum or even none, depending on the use case.

Finally, for gNBs transmitting near prospective lease boundaries, field measurements and power simulations can help estimate to what extent same-frequency neighboring gNBs may interfere. Lessors and lessees must determine acceptable tolerances at boundaries, as otherwise there may may be a large radius around the boundary line where inter-MNO interference prevents transmission.

6) Predicting Future Opportunities of Low Spectrum Utilization: Once a pocket of spectrum is identified, the likelihood of this opportunity recurring is estimated by searching for periodic repetitions (e.g., 1-4am every morning / every Sunday). Customer activity and area may vary, creating size-fluctuating opportunity pockets even within with a general trend of under-utilization. If a weekly pocket occurred in 9 out of the last 10 weeks, there is only a 10% risk it will not recur next week. Lessors may opt for weekly/daily leases if they trust the trend but may not want to offer leases too far into the future, as confidence diminishes over time.

All discussed risks (including if allowing over-QoS-capacity past traffic estimations) are combined to form a final risk for a particular potential lease listing, whose precise formula is left to future work. In general, a one-time listing will be less risky than a repeating listing due to the distance of its ending time to the present, but repeating or larger-area listings may be more profitable and attractive to lessees. The same is true for a listing that offers more coverage or spectrum, which will increase both risk and profit potential. Each lessor must carefully balance risk with potential earnings when choosing which opportunities to list on the Spectrum Marketplace.

7) Differences for Lessees: Prospective lessees can also use this algorithm to find over-capacity regions with some changes. Unlike a lessor, which might conservatively consider a cell under-utilized only when all of its overlapping area polygons are on average under-utilized, a lessee can still benefit even if only part of a cell's coverage area is over-utilized. However, lessees must additionally identify which prospective bands are compatible with its gNB/UE hardware and intended purpose. For example, a commercial MNO may already have base spectrum and want to lease supplementary uplink/downlink using Carrier Aggregation (CA), which may be simpler than setting up entirely new cells if Master/Secondary Information Blocks (MIB/SIB) stay in the base cell and supplementary bands only provide greater capacity / coverage radius. Ample backhaul must also be available to handle increased load.

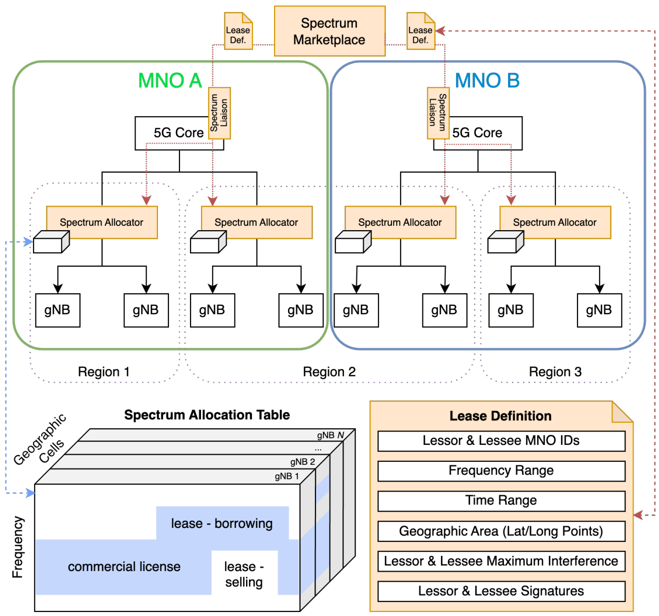

Fig. 7: Spectraleas's Interoperable Spectrum Leasing Architecture

B. Advertising Opportune Lease Listings

Once an MNO determines that an opportunity to lease spectrum discovered by the Spectrum Opportunity algorithm is worth the risk, it constructs a complete lease listing. This listing specifies: 1) a contiguous leased frequency range, 2) lease start/end times (multiple sets for periodic leases), 3) a geographic polygon area for the lease, 4) a maximum tolerable transmit power at the boundaries of said geographic area, 5) pricing information, and 6) lessor/lessee role.

Spectraleas's spectrum leasing architecture is illustrated in Figure 7 which depicts two independent 5G networks connected by a third-party Spectrum Marketplace service that facilitates the transfer of leases. The MNO submits created lease listings to the Spectrum Marketplace which tracks all listings. This is analogous to CBRS's SAS which allocates 10MHz leases to users, but unlike a SAS, the Spectrum Marketplace can *irrevocably* lease any frequency range between any spectrum owners at a fine time granularity, and communicates leases to the RAN. An MNO looking to *obtain* spectrum (as a lessee) can also submit a buy-listing to publicize its interest. In a dynamic marketplace offering hour/minute-duration lease listings, algorithms (not humans) analyze the current listings for matches, via the marketplace or MNO, and algorithms are covered in past research [17], [16], [26].

Listings may not perfectly match, requiring lessors and lessees to negotiate agreeable lease terms (e.g. changing spectrum/time/area) which is primarily done through the Marketplace but may be out-of-band as long as both parties agree to and sign the same lease listing. For example, a lessor should consider the timescale needs of a lessee, as even if the lessor predicts that a lease could be made for many hours, the lessee may only need spectrum for specific, periodic bursts. Lessors and lessees may also need to agree on guard bands if using FDMA/CDMA (Frequency/Code Division Multiple Access) to prevent interference at frequency band bounds.

C. Agreeing on and Preparing For a Lease

Once a lessor MNO and lessee MNO have negotiated a suitable listing match, the lessor agrees to lease its spectrum to the lessee under the listing's specific constraints. The Spectrum Marketplace subsequently creates and sends the lessor and lessee a copy of the same *lease definition*, which defines the set of constraints to be enforced by both parties to execute the lease. These constraints come from the agreed-upon lease listing but additionally contain IDs and cryptographic signatures of each MNO that can be used to verify their agreement to uphold the constraints, as the bottom-right of Figure 7 shows.

Each MNO runs a centralized *Spectrum Liaison* application that interfaces between the MNO's potentially expansive RAN and the Spectrum Leasing Marketplace by sending/modifying lease listings, signing 'agreed' listing deals, and receiving resulting lease definitions. While this design centralizes the Spectrum Marketplace, decentralized systems are possible where MNOs use a blockchain to host lease negotiations/agreements or even directly transacts with other MNOs (to avoid publicizing their interest to lease). But the centralized approach may offer a much wider exposure of prospective lessors and lessees and allow dynamic, short-lived leases.

Both the lessor and lessee MNOs determine which base stations should be reconfigured to comply with the lease. The lessee may use only some available gNBs to transmit based on e.g., cell tower proximity and demand location. An MNO's gNBs on its 'opposite' side of the lease boundary may also need reconfiguration for other ICIC changes as described later.

Once the affected base stations are determined, each MNO's Spectrum Liaison forwards the lease definition to a set of programs known as Spectrum Allocators that manage one or more of these base stations, as Figure 7 illustrates. The Spectrum Allocator is similar to O-RAN's RAN Intelligent Controller (RIC) running an rApp (non-real-time) and tracks allocations over time in a spectrum allocation table, as shown in the bottom-left of Figure 7. Each allocation represents an exclusive, interference-free grant for a gNB to transmit at a given frequency range at a maximum power level. When a lease begins, the Spectrum Allocator reconfigures all affected gNBs to apply the constraints. A new lease definition removes frequency from the allocation table for lessors and adds new allocations for lessees. In Region 2 of Figure 7, each MNO's Spectrum Allocators can reconfigure their applicable gNBs to start/stop using spectrum following a lease definition.

D. Applying a Lease

When an allocation time begins, the Spectrum Allocator communicates with each of its gNB to modify a cell operating in its frequency range with the lease constraints. A gNB may need a new cell or bandwidth part for new spectrum outside an in-use operating band or contiguous frequency range. DUs/RUs with antenna lobe patterns overlapping the boundary may require reducing transmit power. During cell planning, an MNO must configure neighboring base stations to avoid overlapping same-frequency cell interference, often by coordinating TDD/FDD formats and Slot Format Indicator

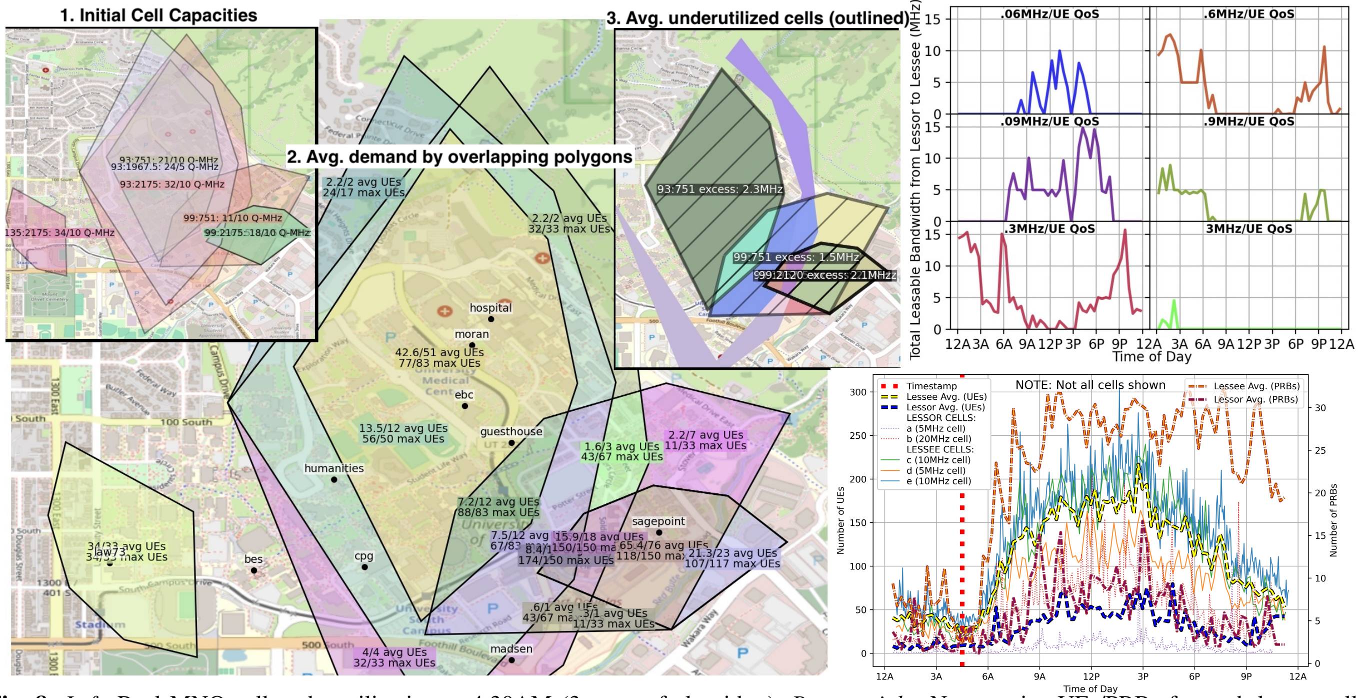

Fig. 8: *Left*: Real MNO cell under-utilization at 4:30AM (3 steps of algorithm); *Bottom-right*: Num. active UEs/PRBs for each lessee cell over one day and lessor/lessee averages; *Top-right*: Total real, beneficial lessor→lessee bandwidth over same day when varying UE QoS

(SFI) patterns. Reconfiguring control channel parameters like UE attach, PCI, PTRS, and SIB broadcast intervals may be needed. Neighbor Cell Relations also require updates for proper handoff, as would Dual Connectivity or Carrier Aggregation setups. The lessor must also relocate all UEs using leased-out spectrum onto alternate cells. At the end of the lease period, lessor and lessee Spectrum Allocator managers revert all RAN configurations.

V. Evaluation

Our prototype implementation of Spectraleas runs the spectrum opportunity algorithm to identify potential leases between MNOs and applies agreed leases. We run the prototype on five outdoor POWDER Ettus USRP B210s base stations with OAI-based 5G Cores and RANs for different mock MNOs, using Quectel RM500Q-GL modules for UEs. All base stations are within one kilometer and use pole-height omnidirectional Commscope VVSSP-360S-F antennas. We use both stationary and mobile UEs (deployed on campus shuttles that traverse the small cells' coverage areas). Servers were d740 nodes.

We developed APIs, data models, and software modules in Python to facilitate spectrum leasing from the inter-MNO agreement phase to gNB frequency configuration. A neutral spectrum marketplace communicates with participating MNOs via Spectrum Liaisons, informing them of new leases. Each MNO runs Spectrum Allocators modules that use lease details to calculate which of their managed gNBs geographically apply to each lease, then reconfigure them as needed (i.e., when a lease starts or ends). Altair Feko is used to simulate gNB power transmission to predict lease boundary interference.

A. Identifying Spectrum Opportunities

To showcase potential spectrum leasing opportunities, we collected real cell measurement data using Falcon real-time monitor software [27] at 10 B210 SDR sites within a 2km area of the University of Utah's campus. Each site decoded LTE Downlink Control Information (DCI) from 50 nearby commercial cells, estimating their UE connection count (via active Radio Network Temporary Identifiers) and PRB load during 1-5 minute intervals. We identified tower cell locations by pairing DCI's PCI (Physical Cell Identifier) data with CellMapper, then estimated sectorized cell coverage areas by angling them towards the midpoint of all sites that detected that cell and fully covering each site, as seen in the left of Figure 8 (measurement sites are shown as black dots, and exact coverage maps would be rougher due to elevation/obstacles).

The bottom-right graph in Figure 8 shows UE load in an area of several all-overlapping lessor MNO cells and lessee MNO cells with combined 25MHz bandwidth/MNO over a 24-hour weekday. By comparing either average UEs or average used PRBs, the lessee generally has much greater traffic per MHz of spectrum, even considering MNOs' diurnal traffic fluctuations. The red vertical line at 4:30AM gives an example when Spectraleas identified areas where 8 lessor MNO cells over a larger area could reduce spectrum usage with minimal QoS impact at 300kbps per-UE QoS. The corresponding left 3 sub-figures show this under-utilization in 3 out of 8 full lessor cells which can be combined in into two 2MHz leasable regions in the 751 and 1967MHz bands. Tripling QoS bandwidth/UE increases leasable bandwidths to 6MHz. Similar patterns occur on later weekdays, decreasing the risk of offering similar early-morning leases in the future.

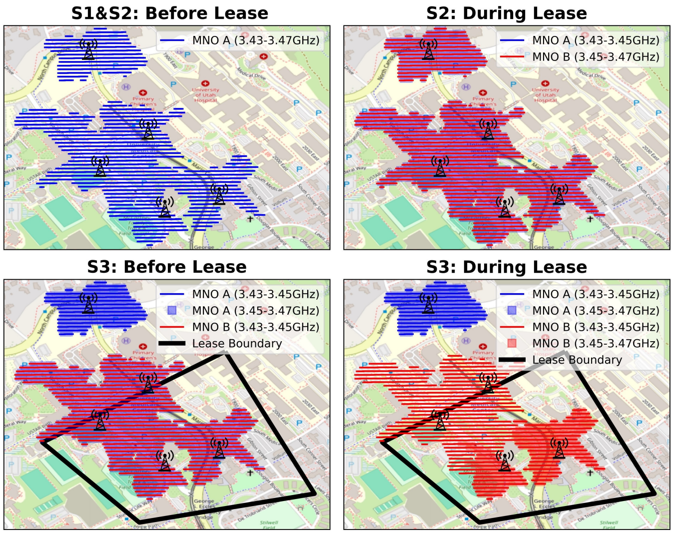

Fig. 9: Three Lease Evaluated Deployment Scenarios

Figure 3 (in Section II) shows these same two leasable regions that overlapped with simultaneous *over*-utilized areas from 4 of the lessee MNO's cells, whose calculations (with the same .3MHz/UE QoS) are omitted for brevity. The spectrum opportunity algorithm estimates that the lessee could have benefited from ~4-8MHz additional spectrum in the 1.35km^2 (334 acres) green overlap area shared with the lessor, meaning 2-4MHz could have been safely leased.

The top-right graphs in Figure 8 depict how total mutually beneficial, leasable spectrum from this lessor to lessee MNO changes depending on time of day and amount of MHz reserved per UE for sufficient QoS. When little MHz is dedicated per UE (.06MHz/UE), there can only be daytime leases where both MNOs had high-enough UE traffic (though the lessee had much more UEs/MHz than the lessor as shown by comparing the difference in their *averaged* traffic in the bottom-right graph). At midnight, this small QoS calculates to *both* MNOs' cells being highly under-utilized, i.e., no 'lessee' demand. In contrast, when much MHz is dedicated per UE (.9MHz/UE) for higher QoS, lease opportunities instead occur in the evening/early morning when few UEs were connected. Outside of the shown .06-3MHz/UE range, the algorithm calculates both MNOs as being either mutually under- or over-utilized, i.e. no supply/demand overlap. In reality, each MNO may value QoS differently, likely between these two extremes.

We have tested the scalability of the spectrum opportunity algorithm's runtime on single measurement periods from 50 to 50,000 cells (on a M1 Max MacBook with 32GB RAM). We consider both sparser and dense scenarios where on average ~5 or ~10 other cells overlap each cell. 5000-cell groups average .40/1.01 seconds (sparser/denser), 50 cells take 0.0021/0.0098 seconds, and 50000 cells take 13.1/38.3 seconds.

B. Spectrum Lease Deployment

After a lease is agreed upon, we run three scenarios (**S1**, **S2**, **S3**) to test the feasibility of Spectraleas's leasing deployment. We measure a lease's impact on the RAN and UEs for both lessor and lessee to demonstrate practical benefits. In all

Scenario	Full (40MHz, A or B)		Half (20MHz, A or B)	
	Tput. (Mbps)	Lat.	Tput. (Mbps)	Lat.
S1	42.2↓ 10.6↑	12ms	N/A	N/A
S2	42.2↓ 10.6↑	12ms	29↓ 8.2↑	13ms
S3 @ X	27.7↓ 10.1↑	13ms	16.8↓ 7.6↑	12ms
S3 @ Y	27.9↓ 9.1↑	11ms	17↓ 7.3↑	13ms
S3 @ Z	14.4↓ 7.2↑	14ms	8.2↓ 6.2↑	14ms
S3 @ Z*	27.8↓ 9.7↑	14ms	17.7↓ 7.2↑	14ms

TABLE I: UE Performance Measurements. *if **not attenuated**

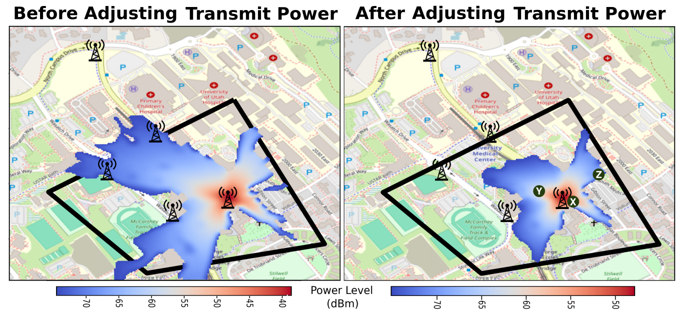

Fig. 10: (S3) Attenuating gNB transmit power to fit within lease. UEs connected to bottom right gNB are measured at X-Z.

scenarios, hypothetical MNOs A (blue) and B (red) operate RAN hardware at five outdoor gNB towers, with A leasing spectrum to B, as shown in Figure 9 (showing cell coverage above -60dBm). In S1, A leases 40MHz to B; in S2, A leases only 20 of its 40MHz to B; and in S3, A leases 20MHz to B only within a boundary around the lower two sites, while B already holds the lower 20MHz license at the lower four sites. Each MNO has two Spectrum Allocator managers that reconfigure 2 and 3 gNBs when the lease starts/ends. UEs are measured 150ft from a gNB with line of sight unless noted.

In **S1**, MNO A leases its entire licensed 40MHz band (3430-3470MHz) to MNO B within an area covering all five sites. After the A-B lease agreement, the Marketplace sends the lease definition to both MNOs' Spectrum Liaisons, which forward it to their Spectrum Allocators. This process took 2.1 seconds on average due to power simulations and area boundary calculation. B's UEs achieve similar performance to A's previous UEs, averaging 42.2 Mbps downlink, 10.6 Mbps uplink, and 12ms latency (see Table I). A must only shut off its gNBs to comply, while B must reconfigure and restart its gNBs, taking 5.8 seconds for OAI. The subsequent UE connection process to B's RAN takes 4 seconds.

S2 introduces more complex, realistic changes in the RAN, such as when a commercial MNO must always provide baseline customer coverage. A leases only the upper 20MHz (3450-3470MHz) to B, retaining the lower 20MHz (3430-3450MHz). A's gNBs initially use the full 40MHz band, so when the lease begins, A's UEs reconnect to the smaller 20MHz band after about 10 seconds of downtime from gNB reconfiguration (gNBs that support dynamic config changes and UE RRC reconfiguration could mitigate this). This results in reduced downlink and uplink throughputs of 29 Mbps and 8.2 Mbps, respectively, as shown in Table I. A and B's UEs experience similar performance since their gNBs use the same hardware, and their gNBs have the same reconfiguration times as in S1.

In **S3**, impactful geographic boundaries are introduced. B already holds the lower 20MHz (3430-3450MHz) at the bottom four gNBs, while A holds the upper 20MHz (3450-3470MHz) across the entire area *and* the lower 20MHz around the top gNB. A agrees to lease its upper 20MHz within a boundary around the bottom two gNBs, but B must attenuate its transmit power by 8.6dBm for the bottom right gNB and 9.7dBm for the bottom left to prevent interference to A (as shown in Figure 10). UE measurements at points X, Y, and Z—different distances from the bottom right gNB—show that attenuated signal strength decreases with distance, leading to lower throughput at point Z, though performance near gNBs (points X/Y) for both 20MHz and 40MHz remains similar to S2 (see Table I). Extra measurements were taken with and without attenuation at farthest point Z, showing that while attenuation decreases speed, the lower signal strength remains useful even near the inner lease boundary edge. Larger inner gaps in coverage are present, highlighting the need to carefully plan boundary coverage for active gNBs. UEs outside the boundary cannot decode attenuated signals, and measurements at A's top gNB are unaffected by the lease. Due to the close proximity of the two center gNBs to the upper boundary, neither MNO can reliably transmit on the leased upper 20MHz in this area (further illustrating potential boundary coverage impacts), though B continues using its lower 20MHz spectrum here without interruption. Power-directionalized sectors and MIMO antennas could mitigate boundary-related interference.

VI. Conclusion

This paper presents Spectraleas, a system that enables short-term dynamic wireless spectrum subleasing. We address the key challenges to lease licensed spectrum temporarily with other MNOs and an algorithm to find leasable spectrum areas amidst varying customer traffic patterns. Furthermore, the Spectrum Marketplace architecture matches prospective spectrum lessors and lessees and applies agreements in the field without disruption to existing spectrum acquisition practices. We evaluate Spectraleas on a two-MNO prototype and demonstrate the feasibility of dynamic spectrum subleasing without negatively impacting customer performance.

Acknowledgements: This work was supported by the NSF through the following grants: #2232463 and #2346555.

References

[1] X. Yang et al., "Mobile Access Bandwidth in Practice: Measurement, Analysis, and Implications," in *Proceedings of the ACM SIGCOMM 2022 Conference*. ACM, 2022, pp. 114–128.

[2] "SMaRT-5G Project - Sustainable Mobile and RAN Transformation." [Online]. Available: https://opennetworking.org/wp-content/uploads/2023/07/ONF-SMaRT-5G-White-Paper-v11.pdf

[3] "Unique spectrum sharing arrangements impacted connectivity during covid-19." [Online]. Available: https://tinyurl.com/mwn7nwn3

[4] J. C. Gallagher, "Repurposing 3.1-3.55 GHz Spectrum: Issues for Congress," 2023, https://crsreports.congress.gov/product/pdf/IF/IF12350.

[5] K. Pelechrinis, P. Krishnamurthy, M. Weiss, and T. Znati, "Cognitive radio networks: realistic or not?" *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 2, p. 44–51, Apr. 2013.

[6] J. Breen et al., "POWDER: Platform for Open Wireless Data-driven Experimental Research," in *14th International Workshop on Wireless Network Testbeds, Experimental Evaluation & Characterization (WiNTECH '20)*, 2020.

[7] M. Dano. How AT&T's network chief hopes to cut a $1.6B electricity bill. LightReading. [Online]. Available: https://www.lightreading.com/sustainability/how-at-t-s-network-chief-hopes-to-cut-a-1-6b-electricity-bill

[8] H. Frank, C. Colman-Meixner, K. D. R. Assis, S. Yan, and D. Simeonidou, "Techno-Economic Analysis of 5G Non-Public Network Architectures," *IEEE Access*, vol. 10, pp. 70 204–70 218, 2022.

[9] M. Wen et al., "Private 5G Networks: Concepts, Architectures, and Research Landscape," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 1, pp. 7–25, 2022.

[10] L. Lopez-Lopez, M. Matinmikko-Blue, M. Cardenas-Juarez, E. Stevens-Navarro, R. Aguilar-Gonzalez, and M. Katz, "Spectrum Challenges for Beyond 5G: The case of Mexico," in *2020 2nd 6G Wireless Summit (6G SUMMIT)*, 2020, pp. 1–5.

[11] X. Yuan, M. Wu, Z. Wang, Y. Zhu, M. Ma, J. Guo, Z.-L. Zhang, and W. Zhu, "Understanding 5g performance for real-world services: a content provider's perspective," in *Proceedings of the ACM SIGCOMM 2022 Conference*, ser. SIGCOMM '22. New York, NY, USA: Association for Computing Machinery, Aug. 2022, p. 101–113.

[12] X. Yang, H. Lin, Z. Li, F. Qian, X. Li, Z. He, X. Wu, X. Wang, Y. Liu, Z. Liao, D. Hu, and T. Xu, "Mobile access bandwidth in practice: Measurement, analysis, and implications," in *Proceedings of the ACM SIGCOMM 2022 Conference*. ACM, 2022, pp. 114–128.

[13] M. Parvini et al., "Spectrum Sharing Schemes From 4G to 5G and Beyond: Protocol Flow, Regulation, Ecosystem, Economic," *IEEE Open Journal of the Communications Society*, 2023.

[14] W. e. a. Ahmad, "5G Technology: Towards Dynamic Spectrum Sharing Using Cognitive Radio Networks," *IEEE Access*, vol. 8, pp. 14 460–14 488, 2020.

[15] M. Matinmikko, H. Okkonen, M. Palola, S. Yrjola, P. Ahokangas, and M. Mustonen, "Spectrum sharing using licensed shared access: the concept and its workflow for lte-advanced networks," *IEEE Wireless Communications*, vol. 21, no. 2, pp. 72–79, April 2014.

[16] M. Gomez, M. Weiss, and P. Krishnamurthy, "Improving liquidity in secondary spectrum markets: Virtualizing spectrum for fungibility," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 2, pp. 252–266, 2019.

[17] M. Mohseni and R. S. Adve, "Inter-Operator Spectrum Sharing Model and Economic Analysis," in *2021 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2021, pp. 185–192.

[18] M. M. Gomez, L. Cui, and M. B. H. Weiss, "Trading Wireless Capacity Through Spectrum Virtualization Using LTE-A," *SSRN Electronic Journal*, 2014.

[19] M. Gomez, "Wireless Network Virtualization as an Enabler for Spectrum Sharing." [Online]. Available: https://d-scholarship.pitt.edu/35988/

[20] Q. Wu, W. Wang, Z. Li, B. Zhou, Y. Huang, and X. Wang, "Spectrum-Chain: A disruptive dynamic spectrum-sharing framework for 6G," *Sci. China Inf. Sci.*, vol. 66, no. 3, p. 130302, 2023.

[21] Z. Li, W. Wang, Q. Wu, and X. Wang, "Multi-Operator Dynamic Spectrum Sharing for Wireless Communications: A Consortium Blockchain Enabled Framework," *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 1, pp. 3–15, 2023.

[22] P. Munoz, O. Adamuz-Hinojosa, J. Navarro-Ortiz, O. Sallent, and J. Perez-Romero, "Radio Access Network Slicing Strategies at Spectrum Planning Level in 5G and Beyond," *IEEE Access*, vol. 8, pp. 79 604–79 618, 2020.

[23] ETSI, "5G; Management and orchestration; Provisioning," ETSI, Technical Specification TS 128 531, May 2022, version 17.3.0 (2022-05).

[24] L. Bonati, M. Polese, S. D'Oro, S. Basagni, and T. Melodia. NeutRAN: An Open RAN Neutral Host Architecture for Zero-Touch RAN and Spectrum Sharing. [Online]. Available: http://arxiv.org/abs/2301.07653

[25] S. Leutenegger, M. Lopez, and J. Edgington, "Str: a simple and efficient algorithm for r-tree packing," in *Proceedings 13th International Conference on Data Engineering*. Birmingham, UK: IEEE Comput. Soc. Press, 1997, p. 497–506. [Online]. Available: http://ieeexplore.ieee.org/document/582015/

[26] Deloitte, "The Impact of Licensed Shared Use of Spectrum: A report for the GSM Association," 2014. [Online]. Available: https://www.gsma.com/spectrum/wp-content/uploads/2014/02/The-Impacts-of-Licensed-Shared-Use-of-Spectrum.-Deloitte.-Feb-20142.pdf

[27] R. Falkenberg and C. Wietfeld, "Falcon: An accurate real-time monitor for client-based mobile network data analytics," in *2019 IEEE Global Communications Conference (GLOBECOM)*, Dec. 2019, p. 1–7.