Interference management for unlicensed users in shared CBRS spectrum

Ghufran Baig
University of Texas at Austin
ghufran@cs.utexas.edu

Ian Kash
University of Illinois at Chicago
iankash@uic.edu

Bozidar Radunovic
Microsoft Research
bozidar@microsoft.com

Thomas Karagiannis
Microsoft Research
thomas.karagiannis@microsoft.com

Lili Qiu
University of Texas at Austin
lili@cs.utexas.edu

ABSTRACT

The citizen broadband radio service (CBRS) is a newly re-purposed spectrum band in 3550-3700 MHz, reclaiming spectrum occasionally used by radars and other incumbents for mobile data communication. It is also a poster child for future LTE-based dynamic spectrum access systems. At present, CBRS does not manage interference from unlicensed LTE users, which we show can be detrimental for its performance. In this paper we develop F-CBRS, a decentralized spectrum interference management system for unlicensed LTE users in the CBRS band. We first look at how much information can each operator be allowed to conceal and how much it has to be mandated (by a regulator) to disclose, and formally prove that the network can achieve fairness only if all operators share fully verifiable information about Access point (AP) locations and user activity. Using this insight we design a channel allocation scheme to efficiently utilize spectrum and incentivise collaboration. This also includes a simple, non-disruptive channel change scheme to frequently and efficiently change channels to accommodate dynamic traffic and environments. Through simulation and testbed evaluation, we show that we increase throughput of more than 90% of the flow by 80%-100% compared to the current CBRS protocol.

CCS CONCEPTS

• **Networks → Wireless access networks**; *Mobile networks*;

ACM Reference Format:
Ghufran Baig, Ian Kash, Bozidar Radunovic, Thomas Karagiannis, and Lili Qiu. 2018. Interference management for unlicensed users in shared CBRS spectrum. In *The 14th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '18), December 4–7, 2018, Heraklion, Greece.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3281411.3281417

1 INTRODUCTION

Motivation: The citizen broadband radio service (CBRS) is an example of dynamic spectrum access in the US, for which the FCC has provided initial approvals and has a goal of the first deployments occurring in 2018. CBRS is designed to reclaim much of the spectrum in 3550-3700 MHz, which is occasionally used by maritime radars on the coast. If successful, it is likely to lead to similar deployments in other bands and locations (*e.g.*, mm wave spectrum, 2.3-2.4 GHz in Europe).

FCC rules [1] mandate use of a database to protect licensed users from unlicensed ones in such deployments. There is no requirement that the databases regulate interference among unlicensed users, although it is permissible to do so. In current designs, they are largely left to coordinate unlicensed users among themselves. This is particularly challenging in case of CBRS as it uses LTE as the radio technology. LTE has been designed with centralized coordination in mind, where networks are run by a single operator who controls resource allocation in the entire network. In absence of interference control, LTE link throughput can be severely reduced, up to 10x, as we show in Section 2.2. MulteFire [2] is an LTE-based technology that has been proposed for unlicensed LTE operation in unlicensed spectrum like CBRS band. It uses Listen-Before-Talk mechanism for fair co-existence with other users in the same spectrum. This proposal has been under discussion for over 4 years now and still there are no products on the market today. Widespread MulteFire deployment will require development of a new wireless ecosystem from scratch and it is still far from obvious if it would ever be widely deployed.

In this paper, we focus on designing a decentralized spectrum interference management system for unlicensed LTE users in the CBRS band. Our goal is to design a system that achieves fair and efficient spectrum use. We ensure that our system is fully compatible with current LTE standards and can be implemented on available LTE hardware, without any changes at the user side, and with only software modifications at the base station. The unique challenges come from a decentralized LTE environment. Unlike conventional cellular systems, CBRS spectrum is managed by multiple network operators and database providers. Our system needs to prevent misbehaviour and provide the right incentives to yield fair and efficient spectrum use.

Challenges and our approach: We propose F-CBRS, a novel spectrum management component that can be deployed on top of a standard spectrum database and LTE base stations. At a high level, F-CBRS proposes a system architecture that provides a way for multiple databases to coordinate information about their unlicensed users and derive a common channel allocation. This poses several challenges.

CoNEXT '18, December 4–7, 2018, Heraklion, Greece

ACM ISBN 978-1-4503-6080-7/18/12…$15.00
https://doi.org/10.1145/3281411.3281417

The first challenge is how much information needs to be shared between different databases and operators, and whether we need to ensure that this information is verifiably correct. All parties (*i.e.*, regulators, database providers and operators) prefer lighter regulation to avoid legal and technical complexities. We look at different options and, perhaps surprisingly, demonstrate that the *only way* to provide fair access to all unlicensed user is to build a system that shares full, verifiable information about access points locations and respective user activity. We show that the system can be arbitrarily unfair if this is not the case. Consequently, we argue that F-CBRS should mandate the full information exchange between an operator and a database provider, and it should be deployed as an extension of the current CBRS framework. Note that CBRS regulations already mandate similar exchange of the information for other types of users, which is certified and thus verifiable.

Next, any change in traffic demand will eventually yield a new channel allocation and will require numerous APs to change their channel. LTE hardware is designed with a fixed operational frequency in mind and we demonstrate that any channel change causes significant disruption. We leverage existing LTE mechanisms to design a simple, non-disruptive scheme for channel change. This allows us to frequently change channels in the network and achieve greater efficiency without affecting users.

Finally, CBRS nodes from the same operator or partner operators can be even more efficient if they operate on the same channel, because they can use time sharing and gain from statistical multiplexing at a much faster time-scale (Section 2.2). However, in order to achieve time sharing, cells have to be in sync (through GPS or IEEE 1588 if indoor) and have to share a central scheduler, which is often not the case, in particular for nodes from different operators. We propose a channel allocation scheme that leverages this observation. The scheme allocates fair fraction of the spectrum to all participants, whether they use time sharing or not. However, once the spectrum is allocated, those that use time sharing can get even more spectrum through statistical multiplexing. Thus, we incentivise – but do not impose – collaboration among operators for more efficient spectrum use.

We evaluate F-CBRS on a testbed prototype that includes the global channel allocation and spectrum reallocation mechanisms on CBRS small cells. We study large scale behaviour in simulations calibrated on our testbed measurements.

In summary, our main contributions are:

- F-CBRS provides a spectrum allocation policy for various players with conflicting incentives in order to come up with an efficient and fair channel allocation for unlicensed users; we demonstrate that such mechanisms need full and verifiable information sharing (Section 4).
- We derive a novel channel assignment scheme that provides a simple non-disruptive channel change mechanism for LTE that allows frequent allocation changes, and a channel allocation algorithm that incentivises participants to leverage statistical multiplexing in time, further improving efficiency (Section 5).
- We implement and evaluate key parts of the system on commodity CBRS small cells and in large scale simulations. We show an improvement of 2x in median user throughput and 1.8x in median page load times for web applications over current CBRS with no spectrum coordination.

CBRS is one of the first implementations of dynamic spectrum access, and, if successful, has a potential to fundamentally change the way we manage spectrum in the future. We believe our findings can help address one of the main technical challenges of future dynamic spectrum access systems.

2 OVERVIEW AND MOTIVATION

In this section, we give an overview of the CBRS regulations and its radio access network.

2.1 CBRS framework

CBRS spectrum consists of bands in the US currently used by military radars, fixed satellite and wireless broadband, comprising 150 MHz of spectrum between 3550 and 3700 MHz. The FCC has recently started a process of repurposing the spectrum and allowing other users to use the same band, provided that they do not interfere with the incumbents [1]. More specifically, FCC has designed three-tiered spectrum access framework where tiers have a descending order of priority access for using the spectrum.

The first tier includes *incumbents* which consists of all users that currently use these frequency bands. The spectrum is available to them whenever and wherever needed.

The second tier are *priority access licensed (PAL)* users. These users purchase short-term licenses for CBRS spectrum use, with 3 years as the maximum initial term. The licenses are sold per *census tract*, which is a geographical area specified by the US government and varies in size but typically includes about 4000 inhabitants. A PAL user can operate on the spectrum in an area if no incumbent is using that spectrum there.

The third tier are *generalized authorized access (GAA)* users. These are the lowest priority users who do not have to pay for the access. However, the spectrum can be accessed by a GAA user only if no incumbent or PAL user is present.

Spectrum management is performed by spectrum access systems (SAS), a set of spectrum databases that coordinate access, by setting the operating channel and power for PAL and GAA access points. Each SAS database has to be certified by the FCC in order to become operational. PAL and GAA users have to register with an SAS database in order to be allowed to use the spectrum. This database then provides these users the operational parameters for their network ensuring that different tiers of priority access are enforced. Currently, there are 7 applicants for database operators in the certification process [3] and it is unlikely that there will be many more given the limited differentiation between the services they provide.

Different SAS databases have to coordinate through a well-defined protocol [4] and exchange information about incumbents and PAL users. In order to provide sufficient protection to incumbents, this information needs to be propagated to all databases within 60 seconds [4]. If this deadline is not met, the database needs to silence all of its client cells to protect the operation of higher priority incumbents. Therefore, the databases vendors maintain their databases in tight synchronization with each other.

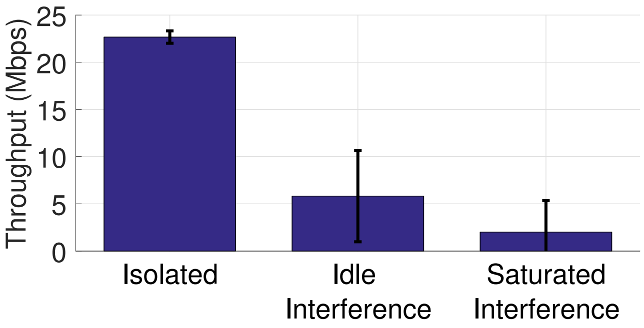

Figure 1: Performance of two non-coordinated and collocated APs. Even when the interferer is idle there is a substantial drop in throughput.

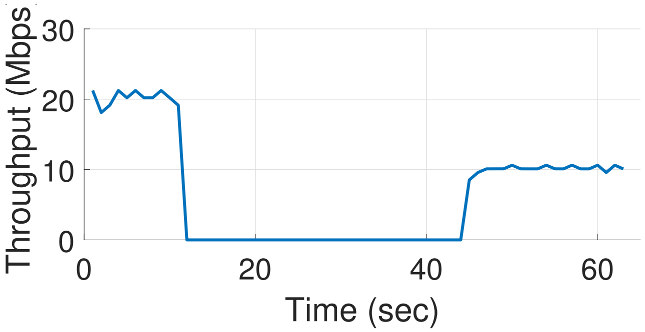

Figure 2: Performance of an LTE client when an AP switches its channel from 10 MHz to 5 MHz.

SAS protocol and the CBRS regulations do not mandate coordination among GAA users and there is currently no requirement to share information about GAA users. It is left to the operators and the database providers to coordinate their deployments with others.

2.2 LTE and coexistence challenges

CBRS is an example of a dynamic spectrum access system that uses LTE radio technology. LTE is one of the most widely used cellular standards today. It is based on a continuous access of an AP and its terminals to its spectrum.

CBRS uses time-division (TDD) LTE technology. The channel is divided into 10ms frames, each further divided in 1ms subframes. Each subframe is further divided in frequency, called a resource block, which carries a data symbol for a particular terminal. A TDD-LTE system shares subframes between uplink and downlink transmissions in one of the preconfigured ratios defined by the standard. This is convenient for dynamic spectrum access as it only requires a single band to operate. However, LTE has two key issues that complicate unplanned deployments: lack of coordination among un-synchronized APs and lack of fast channel switching.

Lack of coordination. In LTE, the ratio and the placement of uplink and downlink slots cannot be configured during system operation. This is in contrast with Wi-Fi where multiple transmissions, uplink and downlink, can be coordinated at a fine time scale using carrier sensing. Therefore, LTE systems have a coordination challenge, and do not have a default provisioning for coexistence. This would lead to transmission collision of uncoordinated LTE links on same channel resulting in huge performance penalties.

To illustrate the co-existence problem of uncoordinated LTE links (such as GAA), we run a simple experiment in which we set up a CBRS AP and connect a mobile terminal to it. We first measure the link throughput in isolation. Then we set up another interfering CBRS AP next to it on the same channel. At first, this interfering AP has no terminals associated to it and it only transmits control signals. Finally, we connect a terminal to the interfering AP and set up a traffic generator that saturates the interfering links. Figure 1 shows that the performance of a link is severely degraded even with an idle interferer, and even more with active interference. Hence, *spectrum utilization can be severely degraded in the absence of coordination*. Section 6.2 further discusses the implications of interference in various setups.

Slow channel switching. LTE networks cannot change the operational channel at fine time scales. They typically operate on a single channel over its lifetime. [1] If an AP stops transmitting control signals, this will disrupt both the control and data channels of the terminal and cause a disconnect.

We illustrate this with an experiment. Since we cannot instantly change the operational frequency of an AP, we set up two APs for the experiment. We set up the first AP on a 10 MHz channel and the second on a different, 5 MHz channel. We quickly decrease the power of the first AP and simultaneously increase the power of the second AP to simulate an event where the AP switches channel. The throughput achieved by the client is depicted in Figure 2. The figure highlights that there is a long period during which the client is disconnected. The disconnection period is lengthy as the terminal is cut off from the network; the terminal needs to perform frequency scanning and search for the LTE synchronization frequency at multiple positions and for multiple channel bandwidths, and subsequently re-attach to the core network. In CBRS, GAA users are required to switch channels as soon as another higher tier user is operational in the area. Moreover, interference coordination among GAA users would also require channel switching quite often. Therefore, such a large switching overhead is unacceptable for a practical architecture for GAA operation in CBRS spectrum.

Synchronization in LTE networks. Centrally orchestrated TDD LTE networks, which we also call *synchronization domains*, can allow for multiple interfering APs to transmit on a single channel. This is achieved by a centralized network controller scheduling traffic across APs for each resource block in every subframe. Multiple synchronized APs can combine their spectrum allocation in a single larger channel, or aggregate their channels using channel bonding to leverage any opportunity for statistical multiplexing. Such networks can synchronize their subframes to sub millisecond accuracy. However, it is unreasonable to expect such a provision to span across all CBRS LTE deployments, especially given the uncoordinated deployment of GAA users; yet, a *synchronization domain* can span networks of a single or a few partnering operators.

---

[1] In case of channel bonding, an LTE AP can change or disable a secondary channel but not the primary one.

Spectrum allocation to GAA users must leverage this LTE feature for efficient spectrum usage.

2.3 Problem definition

As we have seen from previous discussions, CBRS standards do not provision for coexistence among GAA users. GAA users are managed by different entities with conflicting goals (each trying to maximize throughput of their own users). If mutually interfering GAA users are not coordinated, the interference can cause significant disruption and huge performance penalties. For the interfering GAA users not synchronized in time, the only option is to assign them different channels, but frequent channel reassignment potentially induced by varying network load can further reduce efficiency. However, if interfering users are synchronized and coordinated, the system should be able to leverage this coordination to improve spectrum use.

The main questions this paper tries to address are: (a) how to design a spectrum allocation policy: a regulatory framework imposed on GAA users that aligns with incentives of all participants and yields fair access; and (b) how to design an efficient, decentralized network controller that achieves these objectives while being compatible with the existing CBRS standards.

3 F-CBRS ARCHITECTURE

In this section, we give a high-level overview of F-CBRS's architecture. We identify and describe the main building blocks and how they fit in the bigger CBRS architecture and study each block in more detail.

3.1 System Architecture

Owing to strict regulatory requirements, SAS architecture is standardized. Each software component has to undergo an independent certification and all participants need to deploy certified components to make sure protocols requirements are followed. We envisage F-CBRS to become an extension of SAS to deal with GAA users, thus all participants are required to follow F-CBRS protocol and the relevant information reporting. As we discuss later in Section 4, this is the only way to guarantee fair and efficient spectrum distribution.

SAS database providers have a detailed and consistent view of the CBRS spectrum used by incumbents and PAL users. As discussed in Section 2.1, any changes in primaries or PAL users have to propagate to all other databases in 60s, which is a hard deadline. In F-CBRS, we also enforce that all databases have to have a consistent view of GAA users that has to be updated within 60s.

This extension does not impose any significant new synchronization requirement given the small number of expected databases (currently 7). It only slightly increases the amount of traffic to be exchanged from what is currently mandated. As explained next in Section 3.2, we exchange at most 100B of data per AP per 60s interval. A typical census tract with 4,000 users is likely to have less than 1,000 small cells deployed, incurring only 100KB of extra information, about the network, exchanged each 60s. And multiple census tracts can be processed in parallel, this is well within reach with todays reliable datacentre designs (c.f. [5, 6])

We also introduce notion of a *synchronization domain*. As discussed in Section 2.2, operators can synchronize APs within their own network through a controller. Synchronized APs can benefit from statistical multiplexing due to extra coordination in time and resource block scheduling across APs. A synchronization domain is a set of APs that are synchronized and controlled in this way.

We split the CBRS spectrum in 30 channels of 5MHz each. Each AP can be allocated one or more channels. As per LTE standard, it can aggregate any adjacent 5 MHz channels into a single 10, 15 or 20 MHz channel using a single radio [7]. It can further aggregate spectrum using channel bonding.

Standard LTE APs are equipped with a frequency scanner that listens to cell IDs of neighbouring cells and reports back to the operators. F-CBRS requires operators to share this information with the databases so that they can build a global view of the GAA interference graphs. F-CBRS also requires each AP to feature two radios that can simultaneously operate on two different frequencies to implement fast channel switching (discussed in Section 5.1). It is common for today's hardware to dispose of two radio chains, as many small cells are provisioned to operate simultaneously in multiple bands, or feature channel bonding which requires multiple radios. However, having two hardware radio chains is not a strict requirement for F-CBRS. Using radio virtualization [8], these radios can be implemented in software with more complex PHY/MAC chain over a single hardware radio.

An example of a F-CBRS deployment is illustrated in Figure 3 (a). There are two database providers (DB) and three operators (OP), where OP1 and OP2 have a contract with the DB1 and the OP3 has a contract with the DB2. The operators exchange information with their database provider about their networks and APs, depicted with solid arrows. There are two synchronization domains in this example, one comprising AP1 and AP2, and one AP4 and AP5. AP3 and AP6 do not belong to a synchronization domain. The database providers further exchange information about the network among themselves.

3.2 Functional Architecture

CBRS standards [4] dictate that each AP has to report various parameters to its database, including the location, the antenna heights, class, etc. In F-CBRS, we also require each AP to send the following information: (a) The number of active users during the last 60s slot (2 bytes); (b) The identity of the neighbouring APs detected through network scanning and its detected signal strength (*i.e.*, interference graphs, 4 bytes per neighbor); (c) The identity of the synchronization domain it belongs to (if any, 4 bytes per domain). In Section 4 we argue that this is the least amount of information required to achieve fair spectrum allocation. We note that the overhead is low, with at most 100B transmitted per AP during each 60s interval. Note that APs share this information with database providers only. CBRS database service is provided by third party which is trusted by all operators. Database providers make sure that this information is kept confidential and not shared with competing operators due to privacy concerns.

F-CBRS allocates channels in slots where each slot lasts 60 seconds. The intuition for this is three-fold. First, CBRS mandates database synchronization within 60s [4], so we can ensure that

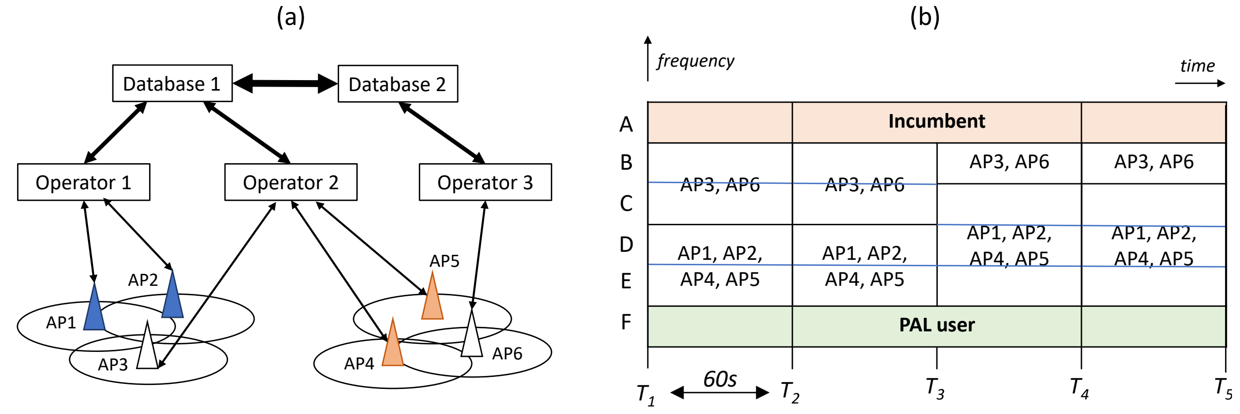

Figure 3: (a) An example of a deployment with two databases, three operators and several Access points, some of which are synchronized; (b) A sample schedule of CBRS spectrum for this network.

databases have a consistent view of the network at the beginning of the slot without extra communication overhead. Second, the LTE connection dynamic has a similar time scale: once an LTE radio sets up a connection, it typically stays connected for 10-20 seconds after sending the last packet due to the data plane setup overhead [9]. Third, the overhead of channel switching has to be significantly lower than the goodput during the interval, and 60s interval meets this requirement.

At the beginning of a slot, each AP reports its information to the database. During the slot, the database exchanges this information along with CBRS mandated parameters with all other databases. Due to CBRS enforced 60s synchronization interval, databases that are unable to sync with the global view silence their client cells for that slot, so all operational databases have the same view of the network at the end of the slot. Further, they all calculate the new allocation, and send it to the corresponding operators and their APs.

The new allocation is calculated from the reported parameters. All databases use the same SINR-based model of the interference that estimates how much throughput a node will get as a function of link length and aggregate interference (see Section 6.2 for relevant measurements). They calculate the aggregate interference on each AP from the interference reports. F-CBRS allows for any fairness metric that can be calculated with the reported information. In Section 5 we define a specific one used in the evaluation.

Because they all have the same view of the network, they are guaranteed to calculate the same allocation by sharing ahead of time any pseudo-random number generator used in the allocation algorithm (described in Section 5). [2] Since PAL licenses are sold per census tract, F-CBRS also derives the spectrum allocation separately and independently for each census tract (noting that F-CBRS can easily be implemented across multiple census tract).

The spectrum allowance per AP is proportional to the number of its active users. This is guided by the notion that the spectrum is shared resource and each user should get the a fair share. This choice is justified in Section 4.

Once the new allocation is calculated, the updated parameters (operating frequency, channel bandwidth and transmit power) are sent to each AP using the standard CBRS messaging protocol. This process requires only a loose time synchronization (100s of millisecond) so NTP is sufficient. If an AP is a part of a synchronization domain then it is also supplied with a list of other frequencies it can use as a part of the domain. The operator's central controller can further adjust frequencies of its APs as long as they don't cause interference to any AP not synchronized with its own.

An example of a channel allocation is given in Figure 3 (b). There are six 5MHz channels denoted with letters A-F. Channel A is allocated to an incumbent, and channel F is allocated to a PAL user. The remaining channels are shared by the 6 GAA users. During the first two slots, T_1 and T_2, AP3 reports the same number of active users as AP1 and AP2 together. Similarly, AP6 reports the same number of active users as AP4 and AP5 together. Consequently, they get the same amount of spectrum: 2 channels for AP3 and AP6, 1 channels for AP1 and AP4, and 1 channel for AP2 and AP5. However, as AP1 and AP2 belong to the same synchronization domain, they can bundle their spectrum into a single 10 MHz channel (D-E); the same holds for with AP4 and AP5. Note that since AP4, AP5 and AP6 do not collocate with AP1, AP2 and AP3, they reuse the same spectrum.

In slots T_3 and T_4, there is an increase in the number of users at AP1, AP2, AP4 and AP5. As a consequence, these APs now get 3 channels: C-E. Since AP1 and AP2 belong to the same synchronization domain, they bundle the 3 channels into one 15 MHz channel and share it in time. The same happens with AP4 and AP5. AP3 and AP6 get one channel: B.

4 SPECTRUM ALLOCATION POLICY

We next discuss policies for allocating spectrum. We have two key desiderata for our allocation. First, we want the allocation to be work conserving: we wish to allocate all spectrum for which there is an interested user. Second, we want the allocation to be fair, the fairness criteria can be arbitrary, depending on what operators want to enforce. To achieve these goals, we adopt the following policy.

---

[2] Alternatively, some database providers may agree to share the calculation of new allocation as well, to minimize computation load.

Spectrum proportional to the total number of active users per AP (F-CBRS): Each operator gets the amount of spectrum on each AP proportional to the number of active users on that AP and inversely proportional to the number of active users on the interfering APs. For allocation to be work conserving, any extra spectrum that can not be used by an interfering AP is also allocated to the APs that can use it. Implementing this policy requires the operators to report detailed information including usage and interference information from each AP (see Section 3 for detailed list of parameters) in a verified fashion (with software certified by a trusted entity, as in SAS database) so that the databases can calculate the correct allocation[3]. Note that this is in line with the current regulators' thinking – for example the FCC certifies CBRS client software to verify the validity of any information it uploads to the database.

In order to justify this choice, we consider some simpler policies below (*CT*, *BS*, *RU*), which require less information to be reported and use an example to show that they can lead to arbitrarily unfair results. We then consider the possibility of operators (unverifiably) self-reporting the number of their active users at each AP and using the same example, we show that this leads to an incentive issue where operators will misreport the locations of their active users. We prove a theorem which shows this problem is inherent to achieving work conservation and fairness in this setting, justifying our choice of policy.

Same spectrum per operator per census tract (*CT*): Each operator gets the same amount of spectrum in a census tract area. This model only requires operators to register with the database.

Same spectrum per AP (*BS*): Each interfering AP gets the same amount of spectrum as others in the vicinity irrespective of the operator it belongs to. Actual amount of spectrum would depends on the number of deployed APs and the interference among them. This model requires AP locations and their interference sensing patterns to be reported to generate interference graph. AP locations are already being shared according to current CBRS SAS rules, and most LTE APs already sense the spectrum for other APs to perform network self-optimization.

Spectrum proportional to the total number of registered users (*RU*):

CT and *BS* assign spectrum irrespective of the number of users which is unfair to users of large operators. Any new entrant to the market can create a huge disturbance in spectrum allocation. RU adds weight to each operator, proportional to its total number of registered customers, allowing operators with more users to get more spectrum. This requires number of registered users to be reported in addition to location and sensing patterns.

All of these policies are clearly work conserving, so to understand their fairness consider the following simple pair of scenarios. There are two census tracts and two operators. Here we assume all APs in one census tract interfere with each other though this is not the case in the general setting. The first operator has n active users at a single AP in the first census tract and none in the second. The second operator has one AP in each census tract. In the first scenario, it has n users in the first census tract and 1 in the second, while in the second scenario it has 1 in the first tract and n in the second. This is illustrated in Table 1.

	Census tract 1		Census tract 2
	Operator 1	Operator 2	Operator 2
Case 1:	n users	n users	1 user
Spectrum:	half	half	full
Case 2:	n users	1 users	n user
Spectrum:	half	half	full

Table 1: Example of unfair allocation (for large n).

All three of these polices (*CT*, *BS*, *RU*) implement a fair outcome in the first scenario (exactly for the first two, and approximately for large n under the third). However, they are all arbitrarily unfair in the second, allocating the second operator (approximately) half the spectrum in the first census tract despite having a single user to the n of the first operator. While this scenario is extreme, it can be thought of as capturing situations where there are "urban" areas where spectrum is in high demand and "rural" areas where it is not, with all three policies advantaging operators which have relatively more rural customers [4].

More generally, we claim that an incentive compatible and work conserving solution is only possible by revealing complete and accurate information, as proposed by F-CBRS. To prove this, we carefully construct a network topology that yields arbitrary bad fairness. Formally, consider a mechanism design problem based on the example from Table 1. Assume all parties know that there are two operators, two census tracts, three APs as above, with the first operator having n_1 active users and the second having n_2. Recall that in this simple example we assume that all APs within one census tract interfere with each other.

We start by considering a set of direct-relevation mechanisms and rely on corollaries from [10] to generalize our observations to any incentive compatible allocation rule. A (direct-revelation) mechanism without payment for this setting is an allocation rule a which takes as arguments the number of active users each operator has in each census tract and returns the amount of spectrum each operator gets in each census tract. That is, $a(x_1, x_2, y_1, y_2)$ gives the allocation where the first operator has x_1 users and the second has x_2 in the first census tract, and the first operator has y_1 users and the second has y_2 in the second census tract. In this example the total number of users $n_1 = x_1 + y_1$ and $n_2 = x_2 + y_2$ is a common knowledge but the operators can choose how many users x_1, x_2, y_1, y_2 they want to report (or misreport) in each census tract.

If a is work conserving then $a(x_1, 0, y_1, n_2)$ results in all the spectrum in the first census tract being assigned to the first operator and $a(0, x_2, n_1, y_2)$ results in all the spectrum in the first census tract being assigned to the second operator. Similarly, if a is work conserving then $a(x_1, n_2, y_1, 0)$ and $(n_1, x_2, 0, y_2)$ result in the operator who did not report all their users as being in the first census tract getting all the spectrum in the second census tract.

Given work conservation, a is fair if $a(x_1, x_2, y_1, y_2)$ results in the first operator getting a $x_1 / (x_1 + x_2)$ fraction of the spectrum in

---

[3] This information is only disclosed to the database providers and is not shared with other operators.

[4] While extending rural coverage has its appeal, this example can cause unfairness in many other scenarios potentially deliberately created by operators.

the first census tract, and a $y_1 / (y_1 + y_2)$ in the second. By the revelation principle [10], it is without loss of generality to only consider incentive-compatible direct-revelation mechanisms. a is incentive compatible if each operator weakly prefers his allocation when he reports the true locations of his users to every other allocation he could get with a different report.

Theorem 1. *Every work-conserving incentive-compatible allocation rule without payment violates fairness. Furthermore, it is arbitrarily unfair for large* n_1*.*

Proof. Suppose a satisfies work conservation and incentive compatibility. As in the example above, consider the two reports $a(n_1, 1, 0, n_2 - 1)$ and $a(n_1, n_1, 0, n_2 - n_1)$. By work conservation, $a(n_1, 1, 0, n_2 - 1)$ and $a(n_1, n_1, 0, n_2 - n_1)$ both assign all spectrum in the second census tract to the second operator. This is because it is known that the first operator has no APs in the second census tract and cannot ask for spectrum.

Thus by incentive compatibility the second operator would report whichever gave him more spectrum in the first census tract, as the spectrum assigned in the second tract for $a(n_1, 1, 0, n_2 - 1)$ and $a(n_1, n_1, 0, n_2 - n_1)$ is the same. Since these two scenarios have different allocations, at least one of them is unfair.

Suppose a assigns a k fraction of the spectrum in the first census tract to the second operator. If the true scenario is $(n_1, 1, 0, n_2 - 1)$, this results in the user of the second operator getting $k/(1 - k)n_1$ times as much spectrum as each user of the first operator. Similarly, if the true scenario is $(n_1, n_1, 0, n_2 - n_1)$, this results in each user of the first operator getting $(1 - k)/k$ times as much spectrum as each user of the second operator. So the unfairness is $\max(k/(1 - k)n_1, (1 - k)/k)$. To minimize this take $k = 1/(\sqrt{n_1} + 1)$, which results in an unfairness of $\sqrt{n_1}$, which is unbounded. □

In summary, our result states that the only way to guarantee a fair spectrum use among GAA users with LTE radios is to mandate all operators to truthfully report (using certified software, much like the rest of the SAS framework) to database providers the number of active users per AP. Any other scheme can be exploited and potentially cause arbitrary unfairness in the network.

Note that our result applies on any policy based on the operators revealing (truthfully or not) their network parameters and databases calculating allocations based on them. It does not apply on schemes that include auctions and payments. However, such schemes are much more complicated to design and have not yet been successfully tested on problems of this scale, so we leave them for future work.

In order to get insight into how different schemes perform on average, we run a simulation with 3 operators and 15 randomly allocated APs and 150 randomly allocated users (for detailed description of simulations, please see Section 6.4). In Figure 4 we show the box plot of throughput achieved by different users when applying different policies from above. As one can see, the more information is disclosed, the more fair the allocation becomes. F-CBRS where most of the information is disclosed increases the throughput of lowest 10th percentile of users to 2.5x, 2.1x and 1.4x over other schemes with varying level of information disclosure, and similarly the median is 2.1x, 1.8x and 1.7x higher.

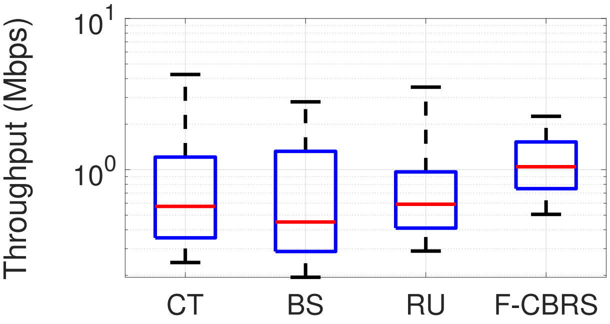

Figure 4: Performance of different schemes

5 CHANNEL ASSIGNMENT

In this section we describe the F-CBRS channel assignment mechanisms. We first describe the fast channel switching mechanism that is required to efficiently deploy the new assignment. We then describe the global algorithm that databases implement to derive the channel allocation for the next slot.

5.1 Fast channel switching

One key requirement to enable F-CBRS's architecture is a fast channel changing scheme for TDD-LTE APs. As discussed in Section 3, an AP may need to change its operating channel every minute. A naive channel change takes almost the same amount of time, as illustrated in Figure 2, thus almost completely disrupting normal operations. Thus, we need to derive a less disruptive channel change mechanism that is *fully backward compatible* with the existing CBRS AP and terminal design.

In order to do this we: (a) impose that each F-CBRS's AP needs to possess two (virtual or physical) [5] radios that can simultaneously operate in different frequencies (as discussed in Section 3), and (b) leverage LTE's standard handover mechanism.

Handover in LTE: Since both virtual radios are co-located they are part of the same Mobility Management Entity (MME). In this setting there are two possible types of handover in LTE.

Handover using S1 interface: In this case the signalling is done through the core network. During the time when handover is in place the packets on data path are either dropped or rerouted through the core network resulting in throughput loss. This throughput loss makes this type of handover very inefficient for frequent channel switching.

Handover using X2 interface: X2 handover is completed without the core network's involvement and happens almost entirely between two APs. Only a single message to the core is required at the end to announce that the handover has taken place. This handover requires direct connectivity (X2 interface) between both APs. In our settings this can be always ensured since both APs are co-located. During the time when handover is in place the packets on data path are also forwarded on X2 interface, hence there is no disruption to the data path.

---

[5] Most LTE base stations today support operation on multiple frequency bands for carrier aggregation

We leverage LTE X2 handover mechanism to achieve fast channel switching on an AP without any throughput loss on data path (as we show in Section 6.3). This is achieved as follows. During the normal operation only one radio operates and the other one is idle. Before the end of each interval, the secondary radio sets itself up in the newly assigned channel and starts transmitting control signals. The primary and secondary APs exchange standard X2 Application Protocol (X2AP) messages between them. At the moment when the channel change is required the primary radio sends handover command to the LTE terminal, which associates itself with the secondary radio. Once the hand-over is performed, we completely switch off the primary radio and make it secondary, while the secondary one becomes primary.

5.2 Global channel allocation

As discussed in Section 3, at the beginning of each slot all databases have a consistent view of the spectrum including the information about the incumbents and PAL users, the number and the location of GAA APs, the number of active users at each one of them and the interference graph describing which AP interferes with which. In this section we describe the algorithm that we run on this data to calculate channel allocation.

The problem is in many ways similar to channel allocation mechanisms applied to Wi-Fi or LTE SON with one main difference: F-CBRS's channel allocation mechanism has to be able to assign non-interfering channels to every AP equal to it's "fair" share, but at the same time prioritize assigning adjacent channels to interfering APs and same channels to non-interfering APs that belong to the same synchronization domain. This is done to maximize the opportunity of aggregating the spectrum allocated to multiple synchronized APs in a single larger band to enable statistical multiplexing by resource block scheduling among them.

We start by taking one of the state-of-the-art channel allocation algorithm for LTE networks called Fermi [11] as a basic building block for our scheme [6], and we modify it to control the channel assignment such that we can maximize channel sharing among synchronized nodes and reduce the overall interference effect in the network by prioritizing same channel assignment to APs in same synchronization domain. For the fairness metric we use weighted max-min fairness, as defined in [11].

Fermi overview: Fermi takes an interference graph provided by F-CBRS's databases. It modifies the graph by adding extra interference edges to create a chordal graph such that it does not contain cycles of size four or more. Fermi first finds the optimal allocation on the modified graph in $O(|V||E|)$, where V and E are the set of vertices and edges in the interference graph, then removes the extra links and assigns spare channels. We restrict the maximal channel share per AP to 40 MHz, given its two radios with a maximum 20 MHz on each [7].

Calculating a chordal graph is a computationally demanding process. However, the interference graph is static and we only recalculate it once a new AP is added. In our implementation of Fermi we make sure that all topology changes propagated through the network are timestamped so that the outcome chordal graph is always the same for all database providers. Also, we make sure that all database providers use the same pseudo-random sequence in their algorithm so that the outcome is the same for everyone.

Algorithm 1: Assignment

```
Input: Clique Tree $\mathcal{T} = \{C_1, \ldots, C_m\}$, Allocation $A_i \forall v_i \in \mathcal{V}$ and $SyncD$
1  Channels assigned to sync domain $d$, $SyncAsg_d \leftarrow \{\}$, $\forall d \in SyncD$
2  Channels assigned to neighbors of v from same sync domain, $NeighAsgn_v \leftarrow \{\}$ , $\forall v \in V$.
3  Channels available for node v, $Avl_v \leftarrow$ All Channels , $\forall v \in V$.
4  Initialize Assignment, $Asgn_v \leftarrow \{\}$
5  $N \leftarrow Root(\mathcal{T})$
6  while $N$ is not NULL do
7      for $v_0 \in N$ do
8          B $\leftarrow GetBlocks(SyncAsgn_d, Avl_{v_0}) : v_0 \in SyncD_d$
9          B $\leftarrow$ B $\bigcup GetAdjacentBlcks(NeighAsgn_{v_0}, Avl_{v_0})$
10         if $A_v <= maxShare$ then
11             $B_1 \leftarrow BlocksOfSize(B, A_v)$
12             $Asgn_{v_0} \leftarrow Asgn_{v_0} \bigcup MinPenalty(B_1)$
13         else
14             $B_1 \leftarrow BlocksOfSize(B, maxShare)$
15             $Asgn_{v_0} \leftarrow Asgn_{v_0} \bigcup MinPenalty(B_1)$
                $B_2 \leftarrow BlocksOfSize(B, A_v - maxShare)$
16             $Asgn_{v_0} \leftarrow Asgn_{v_0} \bigcup MinPenalty(B_2)$
17         end
18         $Avl_{v_0} \leftarrow Avl_{v_0} \backslash Asgn_{v_0}$
19         $rem \leftarrow A_v - |Asgn_{v_0}|$
20         if $rem > 0$ then
21             $Asgn_{v_0} \leftarrow Asgn_{v_0} \bigcup FermiAssign(Avl_v, rem)$
22         end
23         Remove $Asgn_{v_0}$ from interfering nodes channels
            $Avl_v \leftarrow Avl_v \backslash Asgn_{v_0}, \forall v \in C_j : v_0 \in C_j$
24         Add $Asgn_{v_0}$ to sync domain channels
            $SyncAsgn_d \leftarrow SyncAsgn_d \bigcup Asgn_{v_0}$
            $: v_0 \in SyncD_d$
25         Add $Asgn_{v_0}$ to interfering channels in same sync domain
            $NeighAsgn_v \leftarrow NeighAsgn_v \bigcup Asgn_{v_0}, \forall v$, s.t
            $v \in C_j : v_0 \in C_j$ and $v \in SyncD_d$.
26     end
27     $N \leftarrow nextNodeInLevelOrderTraversal(T)$
28 end
```

F-CBRS channel assignment: The key novel addition to F-CBRS is the channel assignment algorithm that greedily packs APs from the same synchronization domain together. This is done using a level order traversal of the clique tree for available chordal graph. Starting from an arbitrary node in the tree, we assign channels to nodes of the interference graph based on their allocation. The assignment process is summarized in Algorithm 1.

We keep track of channels assigned to nodes that belong to the same synchronization domain and for every node v, we also keep track of the assignments for every interfering nodes belonging to the same synchronization domain (line 1-2, 23-24). When assigning to a node v that belongs to synchronization domain d, we do it in two steps. First we get a set of all contiguous blocks of channels that belong to $SyncAsgn_d$ or are adjacent to allocated channel blocks in

---

[6] Our design is tuned to use Fermi but we believe it could be replaced with another resource allocation algorithm and fairness metric.

$NeighAsgn_v$ that are disjoint with existing assignments to other nodes in the same clique (line 8-9). From these channel blocks we pick a block that fully satisfies v's share (or pick a 20 MHz block if $Share_v > 20MHz$) and faces minimum penalty due to transmission in adjacent channels (line 10-17). The penalty is calculated using the model built from measurements shown in Fig 5(b). Prioritizing these channels for allocation maximizes the channel sharing opportunities among the nodes in same synchronization domain while keeping the adjacent band interference low. A sharing opportunity occurs when an AP has channel(s) available adjacent to its own channels that are not used by any interfering APs belonging to some other synchronization domain.

If v's share is not fully met after first round we allocate the remaining share using all the remaining available channels for that node in the second round. Here too we get all the possible channel blocks and chose the one with minimum adjacent channel interference penalty (line 20).

We note that there may be APs that are not allocated any channel even though they have active users due to not enough channels being available in very dense settings. Our scheme allows such APs to use the channels allocated to APs in same synchronization domain to continue their operation (or, if no domain exists, the channel with the least amount of interference).

Finally, we note that there may be APs with no active users. They still need to operate and transmit control signals in case new users want to join. However, even idle, these APs still create significant destructive interference (Section 6.2). Hence in the allocation algorithm we treat them as if they have a single active user.

6 EVALUATION

In this section we evaluate F-CBRS's performance. We start by briefly discussing our testbed. Next, we present our indoor CBRS channel interference, which we use to build the performance model for the large scale simulations. We then present a small-scale testbed evaluation of F-CBRS. Finally, we present a large scale evaluation of the channel allocation algorithm.

6.1 Testbed description

We deploy 2 Juni JLT625 [12] and 2 Baicells mBS1100 [13] LTE small cells operating on CBRS band. All these cells support GPS and IEEE 1588 synchronization, and X2 handovers between cells. However centralized network controller for resource block scheduling is not yet supported [7].We also deploy four CBRS user terminals from same vendors.

We implement F-CBRS's channel allocation algorithm (Section 5) in Python, which can calculate channel allocations in less than 4s, significantly less than the interval limit of 60s.

The APs we use do not allow us direct access to radio, so for executing fast channel change we bundle two APs as one. One AP acts as the primary radio and one as the secondary. We obtain the new channel allocation from our channel allocation algorithm and set up the new centre frequency on the secondary radio and execute the handover by manually adjusting the transmit powers using a tunable resistor to trigger an X2 handover.

To measure the performance of a synchronization domain, we synchronize the APs using GPS and measure performance with and without interference, and use these numbers to infer the performance of time scheduling.

6.2 Channel measurements

We next present throughput and interference measurements. We interpolate the results of these measurement to derive channel link throughput as a function of signal, interference and channel overlap, which we use in our in Section 5 to calculate the allocations and also in large-scale simulation.

Throughput and range: We have measured range, throughput and interference of LTE links in our lab building. We omit detailed range results due to brevity, but we observe that with 20 dBm radios we can establish links of up to 40m on the same floor and up to 35m on the floors above and below.

Overlapping channel interference: We also quantify effects of various types of LTE interference through a set of lab experiments. We set up two unsynchronised APs next to each other. We place two terminals in their vicinity and assign one terminal to each of APs. We measure the downlink throughput of one (terminal, AP) pair and we use the other to generate interference. In the first experiment, the interfering AP is switched off and we measure the throughput of the terminal attached to the other AP. In the second experiment, we switch on the interfering AP but we switch off the terminal that is configured to attach to it. The AP is thus in the idle mode waiting for the connection and it only generates control traffic. In the third experiment we connect the second terminal to the interfering AP and create fully backlogged traffic.

We run two sets of experiments with this setup. In the first set, the two APs share the same 10MHz channel. These results have already been discussed in Section 2 (Figure 1) where we have seen that in-band interference has detrimental effect on an LTE transmission even when the interfering channel is idle.

In the second set of experiments the first AP uses 10 MHz channel and the interfering AP uses an overlapping 5 MHz channel. The results are depicted in Figure 5(a). Again, we observe that the throughput drops significantly even when the channels are partially overlapping and the interfering AP is idle. We thus conclude that non-synchronized APs should not share the same channel under any circumstance, not even when they are idle.

Adjacent channel interference: We next asked how much does the effect of the interference depend on the separation between the channels. We repeat the previous experiment and modify both the gap between two 10 MHz channels and the relative power difference between signal and interference. The results are shown in Figure 5(b). In the most extreme cases an LTE transmission can affect the adjacent channels. This is similar to what is observed in other radio technologies, and matches the performance of LTE transmit filter, which has a 30dB cut-off. We use these measurements in the channel allocation algorithm in Section 5 to avoid packing strongly interfering nodes in adjacent channels.

Benefits of synchronization: Finally, to measure benefits of synchronized scheduling we set up an experiment in which two APs, synchronized through GPS, transmit in the same channel. Link

---

[7] Both vendors informed us that this feature is under development

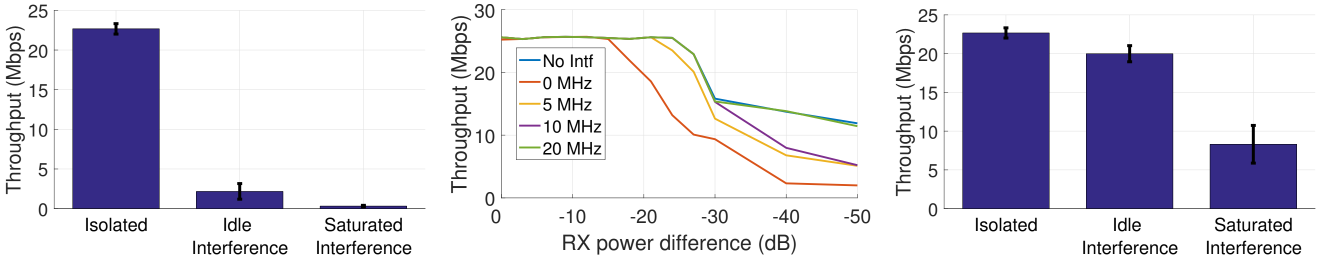

Figure 5: (a) Interference from a partially overlapping channel without synchronization also has detrimental effect on transmission. (b) Throughput in presence of interference from a partially overlapping channel as a function of the gap between channels (0 gap means adjacent channels). (c) Fully synchronized channel, even when fully overlapped, only reduces interference by 10%.

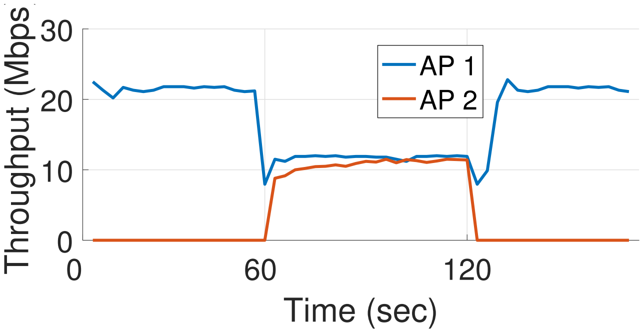

Figure 6: Experiment showing throughput during three intervals. No packet and throughput loss is seen

throughputs are shown in Figure 5(c). Contrary to the asynchronous case, there is very little decrease in throughput when an idle AP coexists in the same channel.

6.3 Testbed evaluation

Here we deploy the end-to-end system on a small test-bed in our lab. We own 4 APs in total. We deploy 2 F-CBRS APs, each consisting of two actual APs, as discussed in Section 6.1. We start with one AP having two users and the other having no users. Then the second AP gets a user. This is fed back to F-CBRS's algorithm which recalculates the shares and conveys them back to the APs.

At the beginning of the new time interval, the APs execute X2 handover and move to the new channels. Note that, due to limitations of our testbed, we perform this operation manually by tuning variable resistors, as discussed in Section 6.1. During the next time interval the user attached to the second AP disconnects. F-CBRS recalculates the share, and at the beginning of the subsequent interval the two APs update their channels. The throughput of the terminals during this experiment is shown in Figure 6. As one can see, the actual throughput closely follows the allocation calculated by F-CBRS's algorithm. We observe no packet losses in the process.

6.4 Large-scale Simulations

We look at the large scale simulations to understand how can F-CBRS affect potential future deployments.

Simulation settings. We implement a link-level network simulator in Python and use measurements from Section 6.2 to derive link-level throughputs.

Topology. We simulate 400 APs and 4000 terminals (corresponding to number of residents in a census tract). We split the APs and terminals across a number of operators (we vary them between 3 and 10, corresponding to the number of currently registered ones). We randomly deploy the network of each operator across the area of interest. Every scenario is repeated 20 times on a new topology. We vary the amount of available CBRS spectrum for GAA users from 100% to 33% (an extreme assuming all of the PAL spectrum is auctioned off) of the 150 MHz CBRS band.

We vary network density by controlling the simulation area. We focus on typical urban area densities [14] as they pose more interference challenges, and we vary from very dense (Manhattan, 70k people per sqm) to sparse (Washington DC, 10k people per sqm). We further assume urban grid model and split the area into buildings of 100m × 100m. Our range measurements to build channel model inside the building, and we add 20dB interference across building [14]. APs and clients are placed randomly within the area.

Workloads. We focus on downlink traffic. We consider two types of traffic workloads. First, backlogged flows for all clients are used for throughput measurements. Second, we model web-like traffic based on realistic parameters regarding flow size, number of objects per page [15] and thinking time distributions [16].

LTE parameters. We use 30dBm transmit power for APs (CBRS category A) and 23dBm for clients (most common chipset limit). Uplink and downlink ratio of TDD LTE is 1:1. Members of the same synchronization domain are able to share channel in time, with the overhead measured in Section 6.2 included.

Spectrum allocation schemes. For reference, we compare F-CBRS with three other spectrum allocation schemes. The first one is a random channel allocation that approximates the current CBRS standards with no spectrum coordination (CBRS). The second one is having operators apply centralized Fermi [11] scheme (Fermi-OP), each on their own network only without considering interference from other operator's network. The third one is having all operators jointly apply centralized Fermi [11] scheme across all CBRS networks (Fermi), corresponds to our scheme without time sharing.

Network throughput. We first look into the link throughput distribution under various spectrum sharing schemes. A bar graph

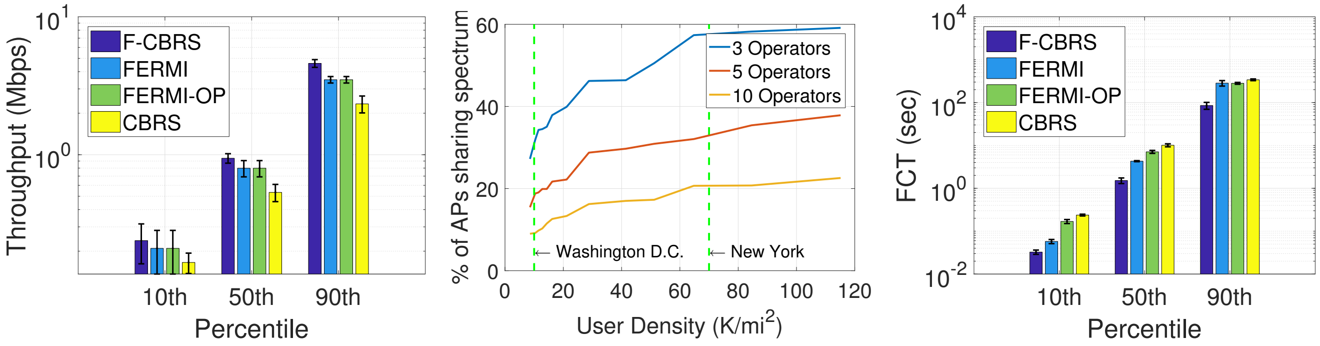

Figure 7: (a) Throughput for large scale simulation experiment comparing 4 schemes showing improvement due to channel sharing. (note the log y-scale) (b) Percentage of APs that get a channel sharing opportunity with another AP in same sync domain without causing interference to any other AP. (c) Page download times across network. (note the log y-scale)

showing average 10th, 50th and 90th percentile of the link rates across the network for the most dense urban area is depicted in Figure 7 (a) (note the logarithmic scale on y-axis). F-CBRS outperforms the next best, centralized Fermi, by 30% in median, 24% for lowest 10th percentile and 27% for 90th percentile, and it outperforms the current (random) CBRS scheme by 2x in median, 1.9x for lowest 10th percentile and 2.2x for 90th percentile. Since all the APs send saturated down-link traffic there is no time-sharing opportunity among synchronized APs, however we still see 30% median improvement over centralized Fermi. This improvement is because F-CBRS prioritize synchronized APs to be on the same channel across the network which have less adverse effect on link throughput, as discussed in Section 6.2. F-CBRS also reduces adjacent channel interference by prioritizing channel blocks adjacent to APs with low RX power.

We repeat the simulations for varying network densities, number of operators and spectrum availability. We observe that the improvement over Fermi decreases to 18% and 14% for median and 10th percentile respectively, and to 75% and 45% for median and 10th percentile respectively over unmanaged CBRS, for a less dense network (10K users per sq. mile) as APs project less interference on others hence reducing the opportunity for improvement for F-CBRS. Also, decreasing spectrum availability reduces the overall network throughput but relative throughput improvement of F-CBRS stays similar. We omit detailed figures for sake of space.

Spectrum sharing. We next look at the effects of sharing and how much sharing opportunity is there. In each of the simulation runs we look at the fraction of the APs that are able to share spectrum in time by F-CBRS. We vary the network density, the number of operators and the available spectrum. This is depicted in Figure 7 (b). When the density is low or the number of operators is high, there is inherently less opportunity for sharing in time as there are fewer nodes in each synchronization domain. As the density increases and the number of operators decreases, the opportunity to share can include as many as 60% of all APs.

Application-level performance. Finally, to understand F-CBRS's impact on real applications, we model dynamic traffic conditions based on our web workload, and examine web-page download times. Here, unlike in the case of fully backlogged traffic, we get the benefits of sharing in time, but there is also less interference with less traffic, so the overall relative improvements are broadly the same. Figure 7 (c) presents the average 10th, 50th and 90th percentiles of page completion times (note the logarithmic scale on y-axis). The reduction in page load times compared to the second best (Fermi) are 40%, 60% and 60% for the 10th, 50th and 90th percentile and 80%, 80%, and 70% for the 10th, 50th and 90th percentile when compared to the current (unmanaged) CBRS scheme.

7 RELATED WORK

Dynamic spectrum access. CBRS spectrum is an example of Dynamic spectrum access (DSA). It has recently emerged as an attractive new innovation opportunity for wireless community with its multi-tiered spectrum sharing model. Several recent papers have presented many design proposals of various parts of CBRS spectrum access architecture. [17, 18] present architectural implementation of a SAS system which mainly focuses on the protocols regarding information sharing between the databases to implement the rules of 3-tiered spectrum sharing for CBRS. [19–23] focus on measurements and studies for coexistence between radar incumbent first tier networks with small cell operators in CBRS band. However spectrum sharing for interference management among GAA users has not been studied in detail. TV whitespace (TVWS) spectrum is another example of DSA which has seen alot of work. TVWS has different set of rules for unlicensed operation and different level of provisioning for incumbent operation compared to CBRS band. The operation of unlicensed users in TVWS is also regulated by databases. Most of the works in TVWS have proposed using WiFi based radio technologies like 802.11af[24], 802.22[25] and WhiteFi[26] for unlicensed operation. Database construction has also been studied for TVWS, [27–29] propose systems for efficient spectrum database construction in TVWS. CellFi[30], proposes an architecture for LTE based unlicensed cellular operation in TVWS which uses distributed interference management using passive sensing to enable spectrum sharing. Similar techniques can be used to minimize intra-channel interference in F-CBRS during the slots once database is done allocating channels to all APs.

LTE for unlicensed spectrum. Several LTE extensions have been proposed that seek to exploit unlicensed spectrum such as LTE-U[31, 32] and LAA[33]. Both these technologies perform energy sensing before starting transmission to ensure medium is shared fairly with other operational devices in the same spectrum. These coexistence mechanisms are only designed to work on the data channel and control channel is assumed to be on a licensed anchor channel operating without any interference. This requirement of a licensed anchor make these technologies unsuitable for GAA operation in CBRS. Moreover, co-existence mechanisms in LTE-U in particular are designed for coexistence with other technologies (e.g. WiFi), and not with other LTE-U devices. MulteFire[2] is proposed LTE small cell technology for standalone operation in unlicensed spectrum that seeks to replace WiFi. It uses WiFi like spectrum sensing and Listen-Before-Talk mechanism for fair co-existence with other users in the same spectrum. Since these features are absent in current LTE standards, MulteFire deployment will require development of a new wireless ecosystem from scratch. This is a major hurdle in MulteFire adoption because of which there are no MulteFire compatible APs or user devices on the market today even when this proposal has been under discussion for several years. Therefore, it is still far from obvious if it would ever be made an industry standard and see wide deployment. F-CBRS proposes an alternative for GAA operation in CBRS that can be deployed using current LTE infrastructure with only software changes at the AP side. Furthermore, all these new LTE proposals do not specify any non-disruptive channel switching mechanism which is required in multi-tier CBRS architecture to cater for higher tier incumbent and PAL users. F-CBRS also benefits from the global view provided by the databases, whereas these LTE proposals are totally oblivious to it.

Channel allocation. Channel allocation for wireless networks is a extensively researched area and there are quiet a few proposals in literature. [34, 35] address resource allocation for WiFi networks, whereas [11, 36] are resource management system designed for OFDMA femtocell networks. We believe F-CBRS can support any channel allocation algorithm or fairness metric proposed in the literature.

8 CONCLUSION

In this paper we have demonstrated that an efficient and fair spectrum allocation policy for various players with conflicting incentives need full and verifiable information sharing among the players, otherwise the spectrum allocation can result in arbitrarily unfairness among them. Using real testbed measurements we have shown that time synchronization among base stations reduces the performance loss due to their interfering transmissions. Keeping these observations in mind we have developed F-CBRS, a spectrum interference management system for unlicensed LTE users in the CBRS band. Our scheme allocates fair fraction of the spectrum to all participants, whether they are synchronized in time or not but incentivises time synchronization by allowing statistical multiplexing to improve the performance of collaborating synchronized nodes. Our large scale simulations show that F-CBRS gives a 100% median throughput increase for users and 60% median performance improvement for web application loads. We believe our insights will be useful in designing future dynamic spectrum access systems.

REFERENCES

[1] FCC. Title 47, Chapter I, Subchapter D, Part 96, Citizen Broadband Radio Service, 2015.

[2] MulteFire release 1.0 technical paper: A new way to wireless. https://www.multefire.org.

[3] Federated Wireless. Google, Federated Wireless, others apply to fill role of SAS, ESC for 3.5 GHz, 2016.

[4] Requirements for Commercial Operation in the U.S. 3550-3700 MHz Citizens Broadband Radio Service Band. Technical report, Wireless Innovation Forum, 2017.

[5] Amazon. Amazon EC2 Service Level Agreement.

[6] Microsoft. SLA for ExpressRoute.

[7] 3GPP TS 36.104v12. Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception, 2014.

[8] Steven S. Hong, Jeffrey Mehlman, and Sachin Katti. Picasso: Flexible rf and spectrum slicing. In *Proceedings of the ACM SIGCOMM 2012 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '12, pages 37–48, New York, NY, USA, 2012. ACM.

[9] Junxian Huang, Feng Qian, Alexandre Gerber, Z. Morley Mao, Subhabrata Sen, and Oliver Spatscheck. A close examination of performance and power characteristics of 4g lte networks. In *Proceedings of the 10th International Conference on Mobile Systems, Applications, and Services*, MobiSys '12, pages 225–238, New York, NY, USA, 2012. ACM.

[10] Roger B. Myerson. Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, February 1981.

[11] M. Arslan, J. Yoon, K. Sundaresan, S. Krishnamurthy, and S. Banerjee. FERMI: A femtocell resource management system for interference mitigation in ofdma networks. In *Mobicom*, 2011.

[12] Juni smallcell & solutions. http://www.juniglobal.com/.

[13] Bai cells. http://www.baicells.com/.

[14] E. Avdic, I. Macaluso, N. Marchetti, and L. Doyle. Census tract license areas: Disincentive for sharing the 3.5ghz band? In *2016 IEEE Global Communications Conference (GLOBECOM)*, pages 1–7, Dec 2016.

[15] J. J. Lee and M. Gupta. A new traffic model for current user web browsing behavior. Technical Report MSR-TR-2014-132, September 2007.

[16] Michael Butkiewicz, Harsha V. Madhyastha, and Vyas Sekar. Understanding website complexity: Measurements, metrics, and implications. In *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '11, pages 313–328, 2011.

[17] C. W. Kim, J. Ryoo, and M. M. Buddhikot. Design and implementation of an end-to-end architecture for 3.5 ghz shared spectrum. In *2015 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 23–34, Sept 2015.

[18] Google. Google Access Services, Spectrum Access System: Managing Three Tiers of Users in the 3550-3700 GHz Band.

[19] A. Khawar, I. Ahmad, and A. I. Sulyman. Spectrum sharing between small cells and satellites: Opportunities and challenges. In *2015 IEEE International Conference on Communication Workshop (ICCW)*, pages 1600–1605, June 2015.

[20] M. Ghorbanzadeh, E. Visotsky, P. Moorut, W. Yang, and C. Clancy. Radar inband and out-of-band interference into lte macro and small cell uplinks in the 3.5 ghz band. In *2015 IEEE Wireless Communications and Networking Conference (WCNC)*, pages 1829–1834, March 2015.

[21] A. Khawar, A. Abdel-Hadi, and T. C. Clancy. Spectrum sharing between s-band radar and lte cellular system: A spatial approach. In *2014 IEEE International Symposium on Dynamic Spectrum Access Networks (DYSPAN)*, pages 7–14, April 2014.

[22] J. H. Reed, A. W. Clegg, A. V. Padaki, T. Yang, R. Nealy, C. Dietrich, C. R. Anderson, and D. M. Mearns. On the co-existence of td-lte and radar over 3.5 ghz band: An experimental study. *IEEE Wireless Communications Letters*, 5(4):368–371, Aug 2016.

[23] S. Kim, J. Choi, and C. B. Dietrich. Coexistence between ofdm and pulsed radars in the 3.5 ghz band with imperfect sensing. In *2016 IEEE Wireless Communications and Networking Conference*, pages 1–6, April 2016.

[24] A. Flores, R. Guerra, E. Knightly, P. Ecclesine, and S. Pandey. IEEE 802.11af: A Standard for TV White Space Spectrum Sharing. In *IEEE Communications Magazine*, 2013.

[25] 802.22. https://en.wikipedia.org/wiki/IEEE-802.22.

[26] Paramvir Bahl, Ranveer Chandra, Thomas Moscibroda, Rohan Murty, and Matt Welsh. White space networking with wi-fi like connectivity. In *ACM SIGCOMM*. Association for Computing Machinery, Inc., August 2009.

[27] D. Gurney, G. Buchwald, L. Ecklund, S. L. Kuffner, and J. Grosspietsch. Geo-location database techniques for incumbent protection in the tv white space. In *2008 3rd IEEE Symposium on New Frontiers in Dynamic Spectrum Access Networks*,

pages 1–9, Oct 2008.

[28] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl. Senseless: A database-driven white spaces network. In *2011 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 10–21, May 2011.

[29] Tan Zhang and Suman Banerjee. V-scope: An opportunistic wardriving approach to augmenting tv whitespace databases. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking*, MobiCom '13, pages 251–254, New York, NY, USA, 2013. ACM.

[30] Ghufran Baig, Dan Alistarh, Thomas Karagiannis, Bozidar Radunovic, Matthew Balkwill, and Lili Qiu. Towards unlicensed cellular networks in tv white spaces. In *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies*, CoNEXT '17, pages 2–14, New York, NY, USA, 2017. ACM.

[31] Qualcomm wants LTE deployed in unlicensed spectrum. http://www.fiercewireless.com/story/qualcomm-wants-lte-deployed-unlicensed-spectrum/2013-11-21.

[32] The unlicensed spectrum usage for future IMT technologies – effecient LTE technologie enables better performance and experience. Huawei.

[33] Ericsson first to give smartphone users indoor boost with license assisted access. http://www.ericsson.com/news/1884277.

[34] Arunesh Mishra, Vladimir Brik, Suman Banerjee, Aravind Srinivasan, and William Arbaugh. A client-driven approach for channel management in wireless lans. In *Proc. of IEEE INFOCOM*, 2006.

[35] Thomas Moscibroda, Ranveer Chandra, Yunnan Wu, Sudipta Sen gupta, Paramvir Bahl, and Yuan Yuan. Load-aware spectrum distribution in wireless lans. In *Proc. of IEEE ICNP*, 2008.

[36] Jongwon Yoon, Mustafa Y. Arslan, Karthikeyan Sundaresan, Srikanth V. Krishnamurthy, and Suman Banerjee. A distributed resource management framework for interference mitigation in ofdma femtocell networks. In *Proc. of ACM MobiHoc*, 2012.