

Towards unlicensed cellular networks in TV white spaces

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

Dan Alistarh
IST Austria
dan.alistarh@ist.ac.at

Thomas Karagiannis
Microsoft Research
thomas.karagiannis@microsoft.com

Bozidar Radunovic
Microsoft Research
bozidar@microsoft.com

Matthew Balkwill
Microsoft Research
a-mabalk@microsoft.com

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

ABSTRACT

In this paper we study network architecture for unlicensed cellular networking for outdoor coverage in TV white spaces. The main technology proposed for TV white spaces is 802.11af, a Wi-Fi variant adapted for TV frequencies. However, 802.11af is originally designed for improved indoor propagation. We show that long links, typical for outdoor use, exacerbate known Wi-Fi issues, such as hidden and exposed terminal, and significantly reduce its efficiency.

Instead, we propose CellFi, an alternative architecture based on LTE. LTE is designed for long-range coverage and throughput efficiency, but it is also designed to operate in tightly controlled and centrally managed networks. CellFi overcomes these problems by designing an LTE-compatible spectrum database component, mandatory for TV white space networking, and introducing an interference management component for distributed coordination. CellFi interference management is compatible with existing LTE mechanisms, requires no explicit communication between base stations, and is more efficient than CSMA for long links.

We evaluate our design through extensive real world evaluation on off-the-shelf LTE equipment and simulations. We show that, compared to 802.11af, it increases coverage by 40% and reduces median flow completion times by 2.3x.

CCS CONCEPTS

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

ACM Reference Format:

Ghufran Baig, Dan Alistarh, Thomas Karagiannis, Bozidar Radunovic, Matthew Balkwill, and Lili Qiu. 2017. Towards unlicensed cellular networks in TV white spaces. In *CoNEXT '17: CoNEXT '17: The 13th International Conference on emerging Networking EXperiments and Technologies, December 12–15, 2017, Incheon, Republic of Korea*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3143361.3143367>

1 INTRODUCTION

The goal of this paper is to design an *unlicensed cellular network*, a network which provides cellular-like experience in unlicensed frequencies – long-range coverage for users with mobile devices, but

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CoNEXT '17, December 12–15, 2017, Incheon, Republic of Korea

© 2017 Association for Computing Machinery.
ACM ISBN 978-1-4503-5422-6/17/12...\$15.00
<https://doi.org/10.1145/3143361.3143367>

at the same time follows the successful model of Wi-Fi’s unplanned deployment. To provide coverage, the network should operate in unlicensed spectrum at low UHF frequencies, taking advantage of the recent availability of frequencies for long-range communications in TV white spaces (TVWS). The network should meet the regulatory requirements for TVWS [1, 2] and allow for the deployment of any number of access points without central coordination but with the ability to control mutual interference.

A natural question to ask is *which wireless technology should an unlicensed cellular network be built upon?* Several standards have been proposed for networking in TV white spaces (802.11af[2], 802.22[3], Weightless[4]). Most have been abandoned and the TV white spaces efforts now mainly focus on 802.11af, which is a modification of Wi-Fi. Wi-Fi appears as a natural fit for an unplanned deployment in TVWS. It inherently allows for network co-existence, and 802.11af amendments allow it to operate in TVWS while maintaining its low silicon design cost by significantly reusing existing Wi-Fi design. Yet, Wi-Fi is originally designed for improved indoor propagation; its PHY is not suited for long-range (Section 3.1) and Wi-Fi’s overheads, such as carrier sense and backoff mechanisms, severely limit its efficiency on long range (Section 3.2).

Beyond Wi-Fi, another obvious candidate for an unlicensed cellular network could be LTE/4G, which presents a well developed ecosystem for cellular networking in licensed frequency. It also supports a large number of spectrum bands, including recently added support for parts of 600 MHz spectrum as a result of FCC incentive auction (which coincides with the TV white space spectrum). It is thus natural to ask how well is LTE-based network suitable for deployment in TV white space spectrum. Surprisingly, this option has not been much explored. Conventional cellular technologies, such as LTE, are efficient and provide long range (Section 3.1). But they are tailored to licensed spectrum where interference is managed either through coordinated control protocols or at deployment phase. LTE has no provisioning to avoid primary spectrum users [1]. Further, it has no mechanism to avoid unplanned interference, which can frequently occur in unlicensed bands. In such cases, current LTE design will lead to significant collisions and performance degradation in TVWS (Section 3.2). Further, conventional cellular networks typically represent expensive deployments based on proprietary hardware and software. Standards and interoperability across protocols and networks are often poorly specified, and in general, “cellular” reflects a complex ecosystem, tightly controlled by providers and hardware vendors. Therefore, LTE in its current form faces major challenges in supporting network co-existence.

In this paper, we propose *CellFi*, a TVWS-compliant cellular network architecture built on top of the LTE stack that addresses

these issues. CellFi leverages LTE PHY-layer advantages to achieve coverage. CellFi extends the LTE stack with two new components, a channel selection and an interference management component. The channel selection component interfaces with a TVWS database and is able to quickly vacate a channel once it has been assigned to a primary user. Interference management across CellFi nodes is based on an algorithm for distributed coordination in LTE that requires no explicit communication or coordination between access points. The interference management component estimates the number of contending nodes using standard LTE radio mechanisms and uses this estimate to calculate its own network share. Each access point then strategically and independently selects one or multiple LTE resource blocks according to this derived share, and relies on the standard LTE scheduler to allocate resource blocks to clients. In essence, CellFi proposes a novel, LTE-compatible and distributed channel allocation mechanism which is based on spectrum reservation and is more suitable for long-range cellular coverage than CSMA.

To demonstrate the feasibility of CellFi and evaluate its performance, we perform several indoor and outdoor experiments as well as large-scale simulations. We find that our system can achieve 1km range while maintaining throughput above 1Mbps, and that it can quickly vacate a channel if a primary user appears. Moreover, we find that our distributed interference mechanism is effective and in most examined cases better than the one in 802.11af.

Our contributions can be summarized as follows.

- We demonstrate the limitations of the existing WiFi and LTE protocols when deployed in the unlicensed cellular scenario. These observations lay the foundation for the design of CellFi.
- We develop CellFi, a long-range, LTE-based, TVWS-compliant cellular architecture. The key component of the architecture is a decentralized interference management that is LTE-compatible and allows for unplanned and uncoordinated LTE deployments.
- We demonstrate the feasibility of CellFi through a series of experiments with off-the-shelf LTE equipment. With the exception of interference management component, CellFi access points have been operational for several months connecting under-privileged users in our local area. Our measurements are further used to guide a large-scale evaluation of CellFi.
- We show that CellFi reduces the number of clients starved due to contention by 70%-90% compared to 802.11af and LTE, and increases coverage by 40% and reduces median flow completion times by 2.3 compared to 802.11af.

To our knowledge, CellFi is the first attempt to provide an unplanned and uncoordinated LTE-based network in low-frequency unlicensed TVWS spectrum. Our vision is that CellFi is the first step towards reducing the cost and improving the quality of cellular networking, while at the same time spurring innovation and new research in a so-far tightly controlled network domain.

2 REQUIREMENTS FOR CELLULAR IN TVWS

Our goal is to design an unlicensed cellular network that comprises the main advantages of the Wi-Fi and cellular. Such a network should thus be characterized by wide coverage and be amenable to

unplanned deployments. We start by examining the requirements of a practical, long-range unlicensed cellular network.

Range. One of the main requirements of an unlicensed cellular network, and its main differentiation against traditional Wi-Fi services is range. The TV white space spectrum promises cell range of 1km and above in unlicensed spectrum, as well as better indoor penetration [2]. We thus require a cell to have a coverage of at least 1km. We also require it to have high per-user throughput of at least 1 Mbps, as promised by universal broadband service in many countries [5].

Database access compliance to unlicensed spectrum. The TV White space is currently available for commercial use in the US, Singapore and the UK, and other countries are working on the relevant regulations. However, rules for accessing TVWS spectrum bands are different than the ones regulating Wi-Fi bands. TVWS spectrum is available to unlicensed devices (secondary users) only in the absence of incumbents (TV and wireless microphones, also called primary users). No device is allowed to access the spectrum before checking spectrum availability in a database [2]. TVWS database compliance is thus an important aspect of the network design.

Unplanned deployment in unlicensed spectrum. To achieve the ease of Wi-Fi deployments, we need support for unplanned, uncoordinated deployments. Since no one owns the spectrum, it is very likely that multiple networks will be deployed in the same area and will need to coexist without mandating central coordination, much like Wi-Fi networks coexist today.

Coexistence between disparate wireless technologies in the same spectrum is a hard challenge, and still not fully solved in practice for many technologies (e.g., Wi-Fi, Zigbee and Bluetooth). In fact, none of the current TVWS standards (802.11af, 802.22 and Weightless) attempts to solve inter-technology coexistence problem. In the same spirit, we will mandate decentralized interference management only among nodes using the same network technology, and assume in the future, either one technology will prevail or that a database will make sure different technologies will occupy different, non-overlapping parts of spectrum.

Furthermore, we require our unlicensed cellular design to work entirely in unlicensed spectrum and not require a licensed spectrum anchor, so that anyone, and not only cellular operators, can deploy it. This is in contrast to current LTE proposals for 5 GHz ISM bands [6–8].

3 EXISTING TECHNOLOGIES & UNLICENSED CELLULAR

We focus on Wi-Fi and LTE radio technologies as most of today's high speed wireless networks, and in particular cellular and TVWS networks, are based on one of these two wireless standards. They incorporate most of the modern PHY and MAC layer design constructs which make them very efficient. Furthermore, well developed designs and accompanying ecosystems make them most viable in business terms – client devices with variants of LTE and Wi-Fi are available today for under \$20.

	PHY (Section 3.1)				MAC (Section 3.2)		
	Design	Freq. chunks	Coding rate	Hybrid ARQ	Access	TX duration	Mode
802.11af	OFDM	6-8 MHz	≥ 0.5	no	CSMA	up to 4ms	uncoordinated
LTE	OFDMA	180 kHz	≥ 0.1	yes	Static	1ms subframes	coordinated

Table 1: Summary of differences between 802.11af and LTE

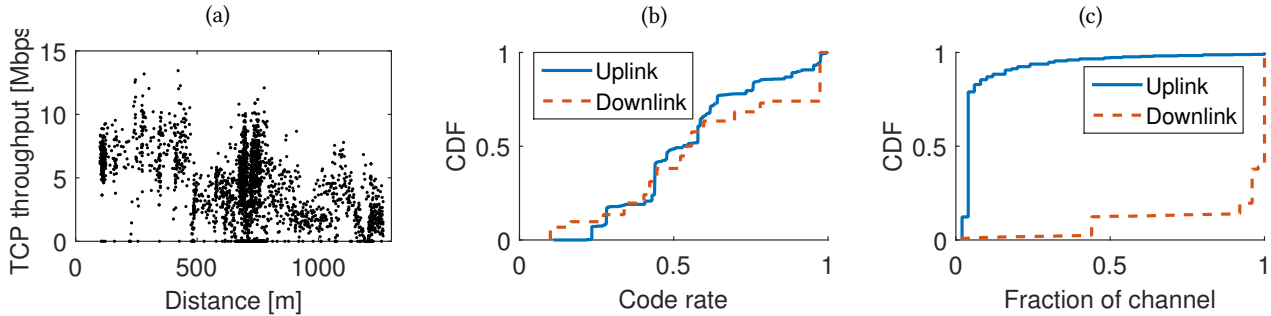


Figure 1: Throughput as a function of distance (a), CDF of coding rate used (b) and fraction of channel used (c).

3.1 Physical layer and coverage

LTE is based on OFDMA. Its clients share common resources, which are defined in frequency and time in terms of *resource blocks* (RB), each 180 kHz wide and 1ms long. Both uplink and downlink resource blocks contain signaling, control and data elements spread out in frequency and time. The signaling elements (reference signals, synchronization signals) are inserted to allow clients to detect LTE transmissions and keep in sync. LTE PHY is designed for long range, so it contains features useful for low SNR links, such as low coding rates, hybrid ARQ and OFDMA modulation that allows the scheduler to use only a subset of resource blocks with the highest signal quality when sending to a specific client.

Most of the LTE hardware already works in a wide range of standardized 3GPP bands [9]. Some of these bands already coincide with TV white space frequencies in parts of the world (e.g., band 44 coincides with part of the TV white space spectrum in the UK). New 3GPP bands are likely to cover even more TV white space spectrum in future (e.g., future LTE bands to cover spectrum sold under US broadcast incentive auctions [10]). 3GPP radio requirements [9, 11] adhere to ETSI TVWS spectral mask requirements [1]. Furthermore, the LTE PHY can use a single channel (in TDD mode) and allows for 5MHz, 10MHz, 15MHz and 20MHz bandwidths; it can thus adapt to several contiguous available TV channels (TV channels are 6MHz in the US and 8MHz in the EU).

Wi-Fi PHY has originally been designed for short range. It uses OFDM, which means that only one client can be served at one time over the entire spectrum, regardless of signal quality on each sub-carrier. Newest Wi-Fi standards (802.11ac) use high coding rates, the minimum being 1/2. The recent 802.11af [2] standard specifies amendments to 802.11 to allow WLAN to operate in the TVWS spectrum. The standard has opted to keep the main features of the 802.11 PHY in order to minimize the cost of modifications. 802.11af PHY uses 6MHz and 8MHz channels and works in TVWS frequencies. It has the same modulation and coding rates as 802.11ac.

Table 1 summarizes the properties of the two technologies. To better understand the impact of these in practice for long-range, we perform an outdoor experiment. We are unable to find 802.11af hardware that operates in the frequencies that we have access to. Instead, we perform an experiment using LTE hardware and monitor the impact of key properties from Table 1, such as OFDMA, coding rate and Hybrid ARQ, on coverage. We deployed an LTE small cell on the top of our building. We moved a client throughout the coverage area and recorded its location, the downlink TCP rate achieved using iperf and various LTE performance metrics (please see Section 6.1 for a detailed description of this experiment).

Figure 1 presents the results of the experiment. We observe that with 36dBm EIRP (29dBm transmit power and 6dBi directional antenna) at the AP and 20dBm transmit power at the client (maximum according to TVWS specs), LTE can reach 1.3km in the urban environment. We measured and achieve 1Mbps TCP rates at more than 85% of measured locations (in Figure 1(a)). In order to achieve these ranges LTE frequently used very low coding rates (Figure 1(b)). In fact, the median coding rate was 1/2, which corresponds to the lowest coding rate offered by 802.11af [2]. Further, LTE leverages its OFDMA capabilities and schedules uplink transmission consisting solely of TCP ACK packages which are small in size in a single resource block. This is shown in Figure 1(c), where we plot the CDF of the fraction of the channel used by transmissions. LTE chooses the resource block with the highest signal strength and improves the quality of transmission – this also explains why the LTE uplink and downlink used similar coding rates. In similar scenarios, WiFi would reduce the signal quality and consequently the range of the network, since it does not implement OFDMA and it would have to send uplink packets across the entire bandwidth. Finally, we observe that LTE leverages hybrid ARQ to improve communication quality, and in particular for longer links - we see that 25% of packets sent from distances larger than 500m use hybrid ARQ.

Overall, we see that the unique features of the LTE physical layer (as shown in Table 1), namely, low coding rates, OFDMA

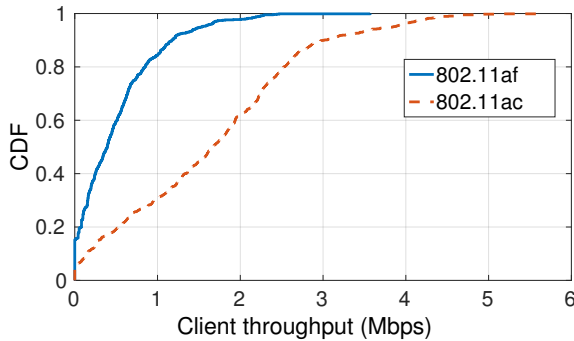


Figure 2: Wi-Fi MAC inefficiencies.

medium access, and hybrid ARQ play a significant role in the LTE link quality and its ability to provide coverage of 1km and beyond.

3.2 Medium access and interference

In LTE, an access point is in charge of scheduling both uplink and downlink traffic. It assigns multiple resource blocks to various clients and the assignment is communicated over the control channel. This makes intra-cell LTE communication very efficient. LTE assumes no unexpected interference from other networks; LTE deployments are well-planned and placement and configuration of access points is such that interference is managed in a coordinated fashion, either from a central network controller or through explicit coordination between neighbouring access points (e.g., through X2 protocols [12]).

In an unlicensed band, these assumptions no longer holds true, as we can witness in numerous unplanned Wi-Fi deployments whose owners make no effort to optimize them. Regulators have also stirred clear from using a TVWS spectrum database to manage interference between unlicensed, secondary users. LTE offers no mechanisms to mitigate interference in uncoordinated deployments, where interference can significantly reduce performance. We illustrate this in detail in an experiment described in Figure 7 in Section 6.3, where we show that a strong interferer (with SINR ≤ 10 dB) can degrade LTE throughput by up to 2 \times , and also cause frequent disconnections. Thus, in order to make unplanned LTE deployment efficient, one needs to manage interference.

The 802.11af [2] standard inherits CSMA medium access mechanism from Wi-Fi, making it suitable for unplanned deployment. CSMA is able to use a channel whenever it is available, and quickly back off and adapt once other users are present.

However, a number of well-known issues in Wi-Fi design render its deployment in TVWS problematic, such as hidden and exposed terminals and scheduling efficiency and fairness [13]. These issues are even more pronounced in a long-range network. To illustrate this, we simulate 802.11af and 802.11ac networks in ns3 (please see Section 6.3.4 for details of simulation settings). In both cases we use 20 MHz channels, and we use RTS/CTS as we have observed that it improves performance. In both cases we consider the same network of access points and place the same number of clients within the corresponding range of each access point. The network range is smaller in case of 802.11ac (home Wi-Fi) than 802.11af (outdoor cellular) because of lower power (20dBm vs 36dBm) and worse

propagation, but the average SNR at the receiver is same in both scenarios. However, the throughput of the 802.11af networks is much worse, as can be seen in Figure 2. Therefore, although CSMA offers efficient and fair contention resolution in shot-range Wi-Fi networks, this is far from obvious in the cellular scenario.

Summary. Overall, LTE appears as a better fit for a TV white space cellular network due to its unique PHY and MAC layer properties. However, it has never been fully considered as a candidate design due to its lack of uncoordinated interference management. We explore this by describing CellFi over the next sections.

4 CELLFI

CellFi enables long-range, unlicensed networking in TV white spaces based on the LTE stack operating as unlicensed secondary users. CellFi incorporates software-based adaptations on the LTE stack to be compliant with requirements of TV white space spectrum access, such as avoiding primary users through a spectrum database, and incorporates decentralized interference management to cater for unplanned deployments. This section provides an overview of the CellFi architecture and its basic building blocks.

4.1 Overview

CellFi is built on the top of standard LTE network architecture. It consists of small cell access points, mobile clients and an LTE control plane (EPC). It also includes a standard TV white space database. Figure 3 presents an overview of the CellFi architecture.

Compared to the traditional LTE, the CellFi access point introduces two new software components, namely interference management and channel selection, and it is equipped with a GPS. The channel selection is responsible for maintaining a list of available channels from a spectrum database and selecting the most appropriate one. The intra-channel interference management component decides which resource blocks within the channel can be used by the access point and which should be left for others, depending on the demand observed for the same channel. The GPS is introduced for two reasons. First, a GPS is required for any TV white space system to provide an accurate location to the spectrum database [2]. Second, CellFi uses TDD LTE allowing it to use a single channel for both downlink and uplink and thus have more flexibility when choosing available spectrum. A GPS clock is indispensable for TDD LTE in order to synchronize interfering uplink and downlink transmissions across multiple networks. Finally, the access point contains the standard LTE small cell software stack that communicates with the control plane (EPC), schedules data packets, etc.

4.2 Channel Selection

Spectrum access in TV white spaces is managed through a spectrum database [14, 15]. The TVWS client in the CellFi access point operates by sending the GPS location details to a TVWS database server [16], to which the database responds with a list of available channels (if any), how long then can be used and maximum allowed power levels. No TVWS client is allowed to transmit in a channel without having a valid lease from a spectrum database and has to stop once a lease has expired.

Satisfying these rules is not straightforward in general. Consider a Wi-Fi client that has associated itself to an AP on a channel with a

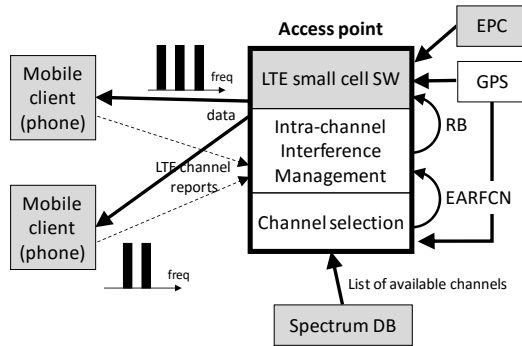


Figure 3: CellFi architecture overview. The shaded blocks denote unmodified existing components (parts of LTE network stack and TVWS database) and the white blocks are new components introduced by CellFi.

valid spectrum lease and has gone to sleep. In order to be compliant, the client needs to verify the validity of the channel before accessing it next time it has a packet to send, which requires changes to the standard Wi-Fi architecture. On the other hand, LTE architecture lends itself to this type of requirements. An LTE client has to get a grant for each uplink transmission from its access point. Thus, once an access point loses a spectrum lease and stops transmitting, all of its clients will stop transmitting instantly (we demonstrate this in practice for CellFi in Section 6.2).

We leverage this observation and build an ETSI-compliant [1] TVWS database client using the PAWS protocol [16]. In our architecture there is a single database client that manages both the access point and all its mobile clients, and all mobile clients have the same generic location parameters [1], determined from the access point’s location. This is because mobile devices may not have a GPS on the device, or may not be able to get the accurate coordinate (e.g., if located indoors).

The access point queries for available spectrum for downlink and uplink independently, and then selects the best TV channel that is available for both downlink and uplink and forwards it to the rest of the LTE stack. It is important to keep in mind that the TV white space database is used only to protect incumbents (TV stations and wireless mics), and not to coordinate spectrum among secondary, TV white space devices. Instead, CellFi uses standard LTE mechanisms such as network listen [17] to find an idle channel from the ones offered by the database, if such exists. If not, CellFi tries to find a channel that is used by other CellFi cells (rather than other non-LTE wireless technologies), as its intra-channel interference mechanism, described next, allows it to gracefully share the channel between other CellFi nodes.

Once a channel is selected, the LTE access points sets the centre frequency (EARFCN) for downlink transmission and announces the uplink frequency in the LTE SIB control message, both in granularity of 100 kHz [18]. LTE clients are required by standards to be frequency agile and to locate all downlink signals within a wide range of frequencies (e.g., existing LTE bands 41-43 are 200 MHz wide), and are allowed to use only the uplink frequency announced in the SIB messages. The database also announces the maximum

transmit powers for the corresponding channels; this also gets communicated to the clients through SIB messages.

4.3 Intra-channel Interference

One of the CellFi design goals is to support co-existence between different networks within the same channel. Conventional LTE access points can coordinate among themselves, using standard protocols (e.g. X2 [17]), to avoid using the same resource blocks for clients that interfere. This however *requires explicit communication and coordination* among access points [19, 20]. In CellFi, coordination is hard to enforce because multiple cellular providers are sharing the spectrum and may not even be aware of one another. CellFi introduces a fully uncoordinated interference management protocol for LTE networks that passively learns about interference from the environment through standard LTE radio procedures, and adapts subchannel allocation accordingly. In this way, much like in Wi-Fi, *no explicit communication or coordination is required*.

The essence of CellFi’s interference management is a short-term resource reservation scheme. Each access point runs a distributed algorithm to decide on a set of resource blocks it will use to serve its users and updates the allocation once every 1 second. It does not have to be explicitly synchronized with others. The intuition for using such long interval (compared to Wi-Fi) is two-fold. Firstly, it amortizes the large channel acquisition overheads that arise in long-range networks (Section 3.2), which make Wi-Fi inefficient. Secondly, such large intervals make sense because LTE also shares a channel in frequency (OFDMA), hence multiple users can be served, each on its own set of resource blocks, during 1-second intervals. In Section 6.3.4 we show that this approach has good efficiency with realistic traffic patterns.

At a high level, the design of the CellFi distributed interference management algorithm can be split into two phases. In the first phase, which we call *distributed share calculation*, each node obtains a conservative estimate of its share of the spectrum, which is roughly based on its share of clients within the neighborhood. In the second *distributed subchannel selection* phase, nodes attempt to converge towards this share by iterating a randomized contention resolution procedure. In this phase, access points attempt to solve an instance of *weighted graph coloring* on the connectivity graph, where their weights correspond to their shares. We discuss these components in more detail in the next section.

Once the interference management component decides which resource block a scheduler can use, it informs the scheduler using standard interfaces. We don’t require any modifications of the standard scheduler. The scheduler is free to schedule any client in any of the resource blocks made available by the interference management system because the interference it creates does not depend on the client selection. This improves spectrum utilization as it allows the scheduler to fully utilize its share of resource blocks by giving all the resources to clients with traffic.

5 INTERFERENCE MANAGEMENT

The interference management component of CellFi needs to solve a distributed channel allocation problem, and in particular it needs to determine: (1) *What share of resource blocks should each network*

get? and (2) Which particular resource blocks should each access point use and how should it adjust it dynamically?

Similar allocation problems have been well studied in other contexts, for example in Wi-Fi or LTE SON. However, there are several specifics that make this problem in the CellFi context unique. Firstly, CellFi is required to manage interference without explicit coordination, unlike conventional LTE networks (c.f. [19, 20]). Secondly, an LTE access point can transmit on several resource blocks at once, and it can change the schedule in each subframe (1ms interval) without any overhead. Further, an LTE client can always sense the status of all resource blocks, even the ones it is not receiving in. This is very different from Wi-Fi where a node can only use and sense one channel at a time, and has significant overhead when changing channels. Thirdly, if two access points transmit on the same resource block and these transmission interfere at a client, the client will not receive its transmission. This is in contrast with Wi-Fi where nodes use CSMA to further avoid interference among nodes that share the same channel. Thus, in comparison with Wi-Fi, CellFi has better sensing and frequency scheduling mechanisms, but the consequences of wrong decisions are more detrimental.

Next, we discuss CellFi's distributed interference management algorithm. CellFi schedules resources in terms of subchannels, where a subchannel is defined to be the minimal set of resource blocks that can be scheduled in LTE and for which we can get channel quality information (Section 5.1). In practice, there are 13 such subchannels on 5MHz channel and 25 subchannels on a 20 MHz channel. The following discussion focuses on the downlink because the uplink is much less saturated; yet, the uplink can be managed similarly.

We start by describing the sensing mechanisms CellFi uses to learn about its neighborhood and then we discuss the distributed share calculation and the distributed subchannel selection phases. We then present the discussion about the convergence properties of the algorithm as well as its theoretical guarantees, showing that the algorithm is guaranteed to converge to the pre-calculated share allocation in $\log(\# \text{ users})$ steps.

5.1 Sensing mechanisms

Like WiFi, the CellFi interference mitigation algorithm relies on sensing information from the environment. The CellFi access point leverages standard LTE radio primitives to estimate the following: **Number of active clients.** In LTE, each client sets up a connection by sending PRACH preambles. This is a special preamble that is used by access points to identify a new node and assign spectral resources to it. In CellFi, we extend this mechanism, and each access point is equipped with an additional PRACH detector that can sense PRACH preambles from *clients it is not serving* (Section 6.3.3). This is used to estimate the number of active clients. The transmit power difference between a client and AP is up to 10dB, and a PRACH detector can reliably detect preambles at -10dB SNR [21]. Thus, any client whose PRACH is detected is likely to be affected by transmissions from the AP. CellFi nodes use PDCCH-order RACH primitive of LTE to solicit PRACH preambles every second [22]. This allows sensing nodes to expire each estimate after 1 second and account for nodes that become inactive.

Client interference in each subchannel. When instructed by its access point, LTE clients report back a channel quality indicator

```

1: function HOPPING( AP  $i$  )
2:    $C_j \leftarrow S_i$  subchannels, picked randomly
3:   for each subchannel  $k$  do
4:     ▸ Draw exponential bucket value
5:      $b_k^i \leftarrow \exp( )$ 
6:   for each phase do
7:     for each occupied subchannel  $k$  do
8:       if  $b_k^i = 0$  then
9:          $k' \leftarrow$  subchannel with maximum utility
10:        swap  $k$  with  $k'$ 

```

Figure 4: Hopping Procedure.

(CQI). The CellFi access point configures its clients to send higher layer-configured aperiodic mode 3-0, sub-band CQI reports [22] every 2 msec. It tracks the maximum reported CQI for each client and *each subchannel* over a period of time. Drops in CQI values indicate interference with a client in that subchannel (Section 6.3.2).

5.2 Distributed Share Calculation

Consider AP i . Let S be the total number of subchannels available, NP_i the number of estimated active clients and N_j the number of active clients associated with AP i . We estimate NP_i using the PRACH detector.

First, for each active client, the AP i reserves S/NP_i distinct shares, giving it a total share of $S_i = N_j * S/NP_i$. This is how we ensure *frequency fair-sharing*. All NP_i clients that AP i interferes with should get enough non-interfering subchannels.

Because of imperfect sensing, this approach can occasionally underestimate the target shares and reduce efficiency, but it is still more efficient than Wi-Fi or LTE, as our evaluation in Section 6.3.4 shows. An AP can also initially overestimate the share available to some of the clients, in which case the scheduler will later automatically assign these to its other clients. This is further discussed in Section 5.4.

5.3 Distributed Subchannel Selection

We now describe the process by which an AP selects and schedules subchannels. For clarity, we split this into the following procedures.

Subchannel Hopping. Initially, AP i randomly picks S_i subchannels. For each subchannel k chosen by i , a random bucket value b_k^i is drawn from an exponential distribution with mean λ (we found $\lambda = 10$ to be a good choice experimentally). Clients associated with an AP send periodic subchannel CQI reports. In all subsequent phases, if a subchannel bucket value b_k^i reaches 0, the AP i gives up subchannel k , and chooses a new subchannel based on a function of CQI values reported by the users that were scheduled on subchannel k . Our implementation chooses the new subchannel that has maximum utility, where utility is defined as the sum of throughput achievable (as estimated from the CQI reading) by all the clients scheduled over the previous subchannel in the recent past scaled by the fraction of time that client was scheduled. See Figure 4 for pseudocode.

Bucket Updates. Each AP updates its bucket values corresponding to employed subchannels periodically as follows. For every client u_j scheduled on the subchannel during the previous period

- If client u_j observes subchannel k as *good* (according to the last CQI report), then b_k^j stays unchanged.
- Otherwise we interpret subchannel k as being a *bad* for AP i . Consequently, the bucket value b_k^i is decremented to $b_k(t+1) = b_k(t) - \text{frac}_j$, where frac_j is the fraction of time that u_j got scheduled on subchannel k during the last period. The bucket update mechanism makes sure that a new AP is able to win a subchannel irrespective of how long the previous AP has been operating on it.

The subchannel hopping and bucket update procedures are similar to other Markovian schemes (e.g. IQ-hopping [23] and references therein) but are adapted to address the main differences between LTE and Wi-Fi, discussed at the beginning of this section.

Channel re-use. Clients very close to their respective access points are not likely to interfere with anyone else; hence, it would be beneficial to schedule them in the same subchannels across different networks to maximize throughput. This is difficult to accomplish without coordination across networks and access points. To achieve this, we use the following heuristic. The access point will give up subchannel i and move to a subchannel of *lower index* if this subchannel is detected as *free* for a certain contiguous period of time, by all of the users that were scheduled on the subchannel i in the recent past. The idea is that clients which experience low interference (such as the ones close to access points), will gradually move towards lower-index subchannels, spontaneously self-organizing. Channel re-use allows for fast convergence and upto 2x gain in throughput for exposed clients as seen in our experiments.




Figure 5: Lines reflect associated clients to access points while dashed lines reflect interference. (a) Information asymmetry with two total channels. eNodeB 1 overestimates his share because it cannot sense UE 2. (b) Information asymmetry with 4 total channels. eNodeB 2 has a share of 1, and while eNodeB 1 can increase his share to 3, it only reserves his fair-share, i.e., 2 channels because it does not know how many subchannels 2 is using (1 due to UEs 3-5)

5.4 Effects of information asymmetry

Wireless sensing depends on node location, and different nodes will obtain different views of the network. Every distributed wireless coordination protocol has to deal with this information asymmetry. The best known examples of information asymmetries in WiFi are exposed and hidden terminals. The distributed share calculation algorithm of CellFi also suffers from two cases of information asymmetry, described below.

Incorrect share: Figure 5(a) shows an example of incorrect share calculation. This case of information asymmetry is dealt with by the scheduler, AP 1 will sense that there are less free subchannels available than it expected, and will not schedule any transmission in subchannels the client is facing interference on, reducing its effective share. The resulting effective share will be feasible, allowing

the distributed hopping algorithm to converge. However, this share adjustment may not be detected by other neighbors of AP 1, leading to possible inefficiencies.

Suboptimal share: Figure 5(b) shows an example of suboptimal share calculation. This case of information asymmetry is fundamental of our setup as AP 1 cannot learn about other clients purely from sensing. It can also not be more aggressive in this case as it could unfairly take a share from AP 2, should the three clients on the right be absent.

The performance effects of information asymmetries of CellFi in general topologies are difficult to analyze precisely. In Section 6.3.4, we show that in complex topologies, CellFi's performance is comparable to state-of-the-art, centralized resource allocation frameworks for cellular networks [20].

5.5 Algorithm Properties

We now analyze the properties of the assignment framework given in Section 5.3. In particular, we will give a sufficient condition under which the basic hopping algorithm (without the channel re-use heuristic) is guaranteed to converge, and probabilistic convergence bounds in this case.

More precisely, we abstract the given setting as an undirected graph $G = (V, E)$, where each vertex $v_i \in V$ corresponds to an AP i . Further, two vertices v_i and v_j are connected by an edge if v_i may interfere with one of v_j 's clients, or vice-versa. Let $N(v_i)$ denote the graph neighborhood of node v_i . Vertices share a set of M subchannels, and initially each vertex v_i has integer demand $d_i \geq 0$, which corresponds to the sum of user shares computed by the algorithm. Our analysis makes two assumptions:

Demand Assumption. In every neighborhood, there exists a *constant factor difference* between the sum of demands and the total number of subchannels.

There exists a constant $1/M < \gamma \leq 1$, such that:

$$\text{for every node } v_i, \sum_{\ell \in N(v_i)} d_\ell \leq (1 - \gamma)M.$$

Fading Assumption. We model subchannel fading by admitting a probability $0 \leq p < 1$ that a subchannel sensed as *free* (and therefore chosen by the hopping procedure) is in fact unusable by the node. This failure event is assumed to be independent of the nodes' random hopping choices, and across hopping rounds.

We focus on *convergence time*, i.e., the time required for the algorithm to reach a configuration in which each node has its subchannel demand fulfilled, and stops hopping.

THEOREM 1. *Under the above assumptions, the algorithm is guaranteed to converge with probability 1. The algorithm will converge in $O(M \log n / ((1 - p) \cdot \gamma))$ rounds, both in expectation and with high probability.*

PROOF. Let us consider the process by which a node v satisfies a unit of its demand. By assumption, the following hold: 1) the node v will not hop on a subchannel currently occupied by another node v' and 2) since a node v_ℓ may only occupy d_ℓ subchannels in a round, and $\sum_{\ell \in N(v)} d_\ell \leq (1 - \gamma)M$, there exist at least $\gamma M \geq 1$ subchannels which are available at every hopping attempt. Therefore, given a

hopping attempt by v , there are two conditions under which it does not succeed in acquiring the channel: either another node makes the same choice (clash), or the channel is faded (fading). We now bound the probability of this event. Both clash and fading events occur independently, with probability at most $1 - \gamma$ and p , respectively. Hence, by assumption, the probability that node v satisfies one unit of demand in a hopping attempt is at least $(1 - p)\gamma$.

Since round choices and fading are assumed to be independent across rounds, the expected time for the fixed node to satisfy a unit of demand is $1/(\gamma(1 - p))$. By a Chernoff bound, for any constant $k \geq 2$, there exists a constant $c \geq 1$ such that the probability that the node fails to satisfy a unit after $k \log n/(\gamma(1 - p))$ consecutive hopping attempts is at most $1/n^c$. By a union bound on the number of nodes, the probability that there exists *some* node which fails after $k \log n/(\gamma(1 - p))$ consecutive hopping attempts, is at most $1/n^{c-1}$. The claim then follows by noticing that a node's demand is of at most M subchannels. \square

It is interesting to consider the effect of channel packing on convergence. Technically, a larger channel slack γ may be needed if hopping and packing occur concurrently, as packing may increase collisions. However, the fact that packing occurs after the node stops hopping ensures that the two procedures are independent to some degree. The empirical evaluation confirms that convergence still occurs with packing, even for dynamic traffic patterns.

6 IMPLEMENTATION AND EVALUATION

We have implemented CellFi in full with the exception of the intra-channel interference management component – current small cell software for the cells we used does not support some of the required standard LTE feature (see Section 6.1 for details). A CellFi access point has currently been operational for several months serving more than 10 users with no broadband connection as identified by a local charity. The need for serving under-privileged population highlights one of the potential use-cases of a long-range, unlicensed, low-cost cellular network. Our users can access the Internet through our gateway using standard LTE clients inside their homes without special outdoor equipment or antennas. In agreement with the requirements identified in Section 2, the network range is around 1km and all users experience rates above 1Mbps. Due to privacy restrictions, we are unable to share any performance numbers related to these users.

Our evaluation covers the main novel components of the system, i.e., channel selection and interference management, through a series of experiments on our testbeds and in simulations. Simulations are used to evaluate the interference management component which we are unable to implement – yet, simulation parameters such as potential inaccuracies of our sensing mechanisms, or interference due to control signaling are guided by our testbed measurements. To this end, we are confident about the realism of our simulated results.

6.1 Implementation Details

We implement our architecture using IP Access E40 small cells [24]. The small cell operates on 3GPP band 13, which we can use in our area. The transmit power of the small cell is 23 dBm. We further

use Amphenol directional antenna with 7dBi gain and about 120 degree sector width.

Channel selection is implemented on a PC. We interface and test it with a certified Nominet spectrum database [25]. We are unable to implement the interference management controller on our current small cells due to lack of software support for the X2 interface and CQI aperiodic mode 3-0 (sub-band level report). Instead, we implement the full component within ns-3 simulator and evaluate its basic mechanisms in a test-bed.

The mobile client used in the experiment is based on a Qualcomm MDM9625 chipset. The client's transmit power is limited to 20dBm, as per TVWS specifications. We use QXDM to get CQI information and other internal signaling information from the mobile client.

We have also built a custom access point on a software designed radio (SDR) that supports a limited set of LTE features; our SDR access point is fully LTE compliant at the PHY level, which we have verified using commercial LTE test equipment. We use the custom access point to introduce controlled interference and evaluate the complexity of the PRACH detector; we are otherwise unable to achieve these using the commercial small cell.

Measurements without SDR are performed outdoors. Measurements with SDR are performed indoors as the SDR was not equipped with the adequate power amplifier to reach the same range.

6.2 Channel selection

We first evaluate our channel selection component by measuring the time it takes for our system to vacate and reacquire a channel, following a change in a database. The experiment is illustrated in Figure 6. We show that the response time is in compliance with ETSI specifications [1] which mandate that transmissions should stop within one minute after the channel ceases to be available.

In general, the granularity of channel availability is expected to be in hours and days [2], as the channel is allocated to the incumbents such as wireless microphones for special events. We confirm this by inspecting the content of the database during the last year.




Figure 6: Spectrum database interaction experiment: At 57sec channel is removed from the DB for 5 min, 2 sec later the AP radio is turned off and the client stops transmitting.

AP takes 1min 36sec to reboot and start the radio after channel is reintroduced. This is because our interface with the AP requires an AP reboot after any radio parameter changes. We note that this process can be shortened using more advanced AP control interfaces such as TR-069 [26].

Once the AP is on, it takes another 56s for a client to connect and resume traffic because it has to perform cell search on various frequencies in multiple LTE bands. This time can be further reduced by disabling unused LTE bands, which we are unable to do on our existing clients.

6.3 Interference management

As discussed in Section 6.1, we are unable to implement the interference management component due to software limitations of our small cell. Instead, we evaluate its key building blocks in a test-bed, and its network-level performance in large-scale simulations.

6.3.1 Interference management with subchannels. Interfering cells can mitigate inter-cell interference by avoiding to use the same subchannels at the same time. This is not immediately obvious as LTE control elements are always present and can create interference even when there is no data being transmitted. This is particularly pertinent when a client is closer to an interfering cell than its serving cell.

To understand the impact of control channel interference, we perform outdoor experiments with two E40 small cells, one acting as a serving cell and the other as an interfering one. The deployment is depicted in Figure 7(a). Both cells are placed on the rooftop of our building, 15m high. The dashed lines depict the direction of each of the two antennas, and the solid blue line depicts the path over which we walked with a mobile device and performed the measurements. The end of the path is at about 250m from our building. We note that we have walked even further but due to the topology of the area the signal remains stronger than the interference and the results were similar to the ones at the end of the current path. We observe a large variability in SINR values, from -15dB to +30dB. We get such low SINR values because one end of the path is in the direction of the interference and outside of the main direction of the serving cell antenna.

We measure performance using Qualcomm’s QXDM tool at the client and we log the received signal levels (RSSI) from one or both cells and the client’s goodput in bits per symbol, where $\text{bit/symbol} = \text{coding rate} \times (1 - \text{BLER})$. We express goodput in bits per symbol rather than application-level throughput because our cell serves other users as well; hence, we measure throughput only within the resource blocks allocated to the test client.

We perform three measurements on the client’s performance on the path: i) when only the serving cell is active, ii) when the interfering cell is active but has no users and iii) with the interfering cell being fully backlogged downstream. Figure 7(b) shows the goodput as a function of the received signal strength for the case of no interference and signalling interference. The two vary by at most 20% and in most cases much less than that. Hence, the control channel interference on its own does not affect the performance of data transmission significantly, even with SINR as low as -15dB. Figure 7(c) presents the CDFs of the achieved goodputs with control channel interference only and with full data interference. We consider only the points where SINR is below 10dB, as in the other cases the goodput does not get affected much. We see that data interference can reduce the throughput by as much as 50% in some cases. Also, we observe frequent disconnects at one end of the path when data interference is present, which we don’t observe

with control channel interference (disconnections are not included in the figure as we cannot register goodput during these intervals).

Overall, data interference is critical in an LTE system. Yet, two cells can share the spectrum successfully if they can coordinate data access in a way that we propose in Section 4.3, despite interference from the control plane. We use these measurement results in our simulations to account for the effects of control-channel interference.

6.3.2 CQI & channel quality. The distributed subchannel selection algorithm is based on detecting subchannel interference from CQI reports (Section 5.1). We now demonstrate that CQI is a sufficiently accurate estimator of interference.

Our CQI estimator needs to balance two challenges. First, channel quality fluctuates due to changes in the environment. Second, an interfering signal might be weakened due to fading and not affect the overall throughput. The estimator should not trigger subchannel reallocation due to mis-identification of interference or when the interference signal is weak as this could result in loss of throughput; this will not allow the network to converge. These effects are highlighted through real measurements in Figure 8. Due to fluctuating channel conditions, throughput varies significantly in the second OFF period, even when no interference is present. Further, the last ON period shows the effect of fading, where despite interference being present, its signal is weak thus not affecting the overall throughput.

While CellFi requires subband CQI reporting, our commercial small-cell access point does not implement the full LTE spec and only reports wide-band CQI over the entire 5 MHz channel. Thus, we only use wideband CQI reporting for this experiment. Yet, the same observations apply.

Our estimator works as follows. Since interference is typically bursty, we consider the maximum CQI observed within a time window as an estimate of CQI for a channel without interference. We declare that interference is present if we observe a CQI report below 60% of this maximum value over a window of 10 consecutive samples. We measure the false positives by running the detector over samples of the channel without interference. We observe that it has less than 2% false positives, i.e., one false positive every 100ms on average (CQI is sampled every 2 ms). Our measurements further show that when interference is strong, our detector correctly reports interference with 80% probability. As with the interference measurements, we use these results in our simulations to model imperfect interference detection.

6.3.3 PRACH preamble detection. CellFi uses PRACH to estimate the number of contending clients. It is known that PRACH preambles can be detected at -10dB [21]. Here, we show that we can design a low-complexity PRACH detector for that purpose.

The key challenge for an access point trying to overhear PRACH preambles from clients not associated with it, is detecting these preambles efficiently without knowing the preamble sequence number [18] or having the timing information. A naive implementation would correlate several long PRACH sequences, one for each preamble sequence number, whenever new samples are received.

We propose a different detector which leverages the structure of the preambles. If we correlate a PRACH preamble received with a time offset, it will have a cyclic shift in its frequency representation.




Figure 7: Outdoor LTE interference experiment setup (a), the difference between no interference and signalling-only interference (b) and the difference between signalling-only interference and data interference (c).




Figure 8: PHY throughput and channel quality indicator (CQI) reported during four states of interfering radio. ON denotes times where there is interference.

Equally, a different cyclic shift will be caused by a different preamble sequence number. We only need to detect whether a preamble is present, and neither need to know which preamble sequence has been transmitted nor when the transmission has started. Thus, we only need to perform two correlations; one to detect the most likely cyclic shift and another to check its correlation value.

We implemented the modified detector on the SDR platform. The detector has a comparable performance to a conventional implementation (with timing information) when receiving PRACH signals from a commercial dongle. Overall, it is 16 times faster than the required line rate (when ran on an Intel i7 CPU on a 10MHz channel).

6.3.4 Large-scale Evaluation. We evaluate the performance and properties of CellFi’s interference management in large scale networks using the ns3 simulator [27]. Simulations are parametrized based on the experiments in the previous section. We model the control channel interference by scaling down the measured throughput based on the measurements in Fig. 7. We have introduced 2% false positives and 80% probability of correct interference detection using the measurements in previous section to model imperfect interference detection because of incorrect CQI reporting. Our measurements show that our PRACH detector can detect the preambles reliably at -10dB, we count only the users whose PRACH can be heard at -10dB, this simulates the imperfect user detection as users whose PRACH is heard below -10dB are still affected by the interference but are not considered in share calculation. Overall, the evaluation examines a set of static and dynamic traffic scenarios

along three dimensions.

- *MAC effects on throughput and coverage* in presence of interference, and comparison against LTE and Wi-Fi.
- *Application-level performance* by measuring the page download times of a web-like workload.
- *Distributed subchannel selection.* We evaluate CellFi’s resource allocation against a centralized, oracle-based state-of-the-art OFDMA resource isolation scheme [20].

Simulation settings. We simulate an area of 2 km x 2km, with a varying network density as controlled by the number of simulated APs. Base stations are randomly placed in this area with varying number of clients per AP. Unless otherwise noted, every scenario is repeated 20 times on a new topology.

Workloads. We consider two types of traffic workloads and focus on downlink traffic. 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 and object size from [28] using thinking time distributions [29] to get flow inter arrival times.

Wi-Fi parameters. We simulate 802.11af by adjusting the standard 802.11ac PHY and MAC layer in ns3 to match the 802.11af specs [2]. Our Wi-Fi implementation uses ideal rate adaptation based on the receiver’s SINR value, and supports MPDU aggregation with maximum possible aggregated frame size of 65 KB. RTS/CTS is enabled; its overhead is small due to the large aggregation and Wi-Fi performance is better with RTS/CTS. The channel bandwidth for Wi-Fi is 6 MHz.

LTE parameters. We use the standard ns3 LTE implementation. For CellFi, we added control channel interference and CQI detection probabilities, both derived from our measurements (Section 6.3). We choose 5MHz channel and TDD type 2, configuration 4 [30] which grants 7 downlink (7ms) and 2 uplink (2ms) subframes in every 10ms frame.

RF. Wi-Fi and LTE support different PHY data rates, yielding different coverages. We choose the transmit powers for the two networks that provide the same coverage in isolation, in order to focus on MAC-layer efficiency. For WiFi, TX power of both AP and client is set to be 30 dBm. For LTE, TX power of AP is 30 dBm and client is 20 dBm. We model loss propagation and noise floor based on our range measurements (Section 3.1).




Figure 9: Coverage vs density for CellFi, Wi-Fi and LTE for 6 clients (a), client throughput of CellFi, Wi-Fi, LTE and the oracle (b), page load times of CellFi, LTE and Wi-Fi (c).

Coverage and throughput. We have already shown in Section 3.1 that the link range can go beyond 1km when there is no interference. Here, we examine coverage when interfering cells are present. We use the static traffic workload and we compare CellFi against Wi-Fi (802.11af) and LTE. We also examine the performance a centralized oracle subchannel allocation [20] can achieve. This provides the extent of performance loss compared to an optimal offline allocation.

Figure 9(a) shows how coverage varies as the network densifies. CellFi improves coverage (number of connected users) and reduces the number of starved nodes compared to both CellFi and LTE. For 6 clients per AP and 14 APs, coverage increases by 37% and 16% compared to Wi-Fi and LTE respectively. In an even denser scenario with 16 clients (not shown due to lack of space), CellFi still offers coverage to more than 80% of users, an increase of 32% and 8% compared to Wi-Fi and LTE.

We now drill down in more detail and examine the performance offered in the densest of the scenarios with 6 clients of Figure 9(a). This reflects 84 concurrent clients in a 5MHz channel; note that a typical 3G macro-cell today supports 32 active clients on the same bandwidth [31]. Figure 9(b) presents the CDF of client throughput achieved across 20 runs of the experiment. We observe that CellFi improves the overall coverage and fairness, without sacrificing the total throughput in the network. On the contrary, while with Wi-Fi and LTE some clients enjoy higher throughput, this is at the expense of 30-40% of starved clients due to exposed and hidden terminals in the case of Wi-Fi, and lack of interference management in the case of LTE. Overall, CellFi roughly doubles the total throughput achieved per AP compared to Wi-Fi at the median, while reducing starved clients by roughly 70% compared to LTE and Wi-Fi. We also observe that CellFi always provides connectivity to more than 90% of the clients, while there are cases were Wi-Fi and LTE only provide connectivity to 30% and 60% of clients respectively. Note that CellFi presents near optimal performance when compared to the oracle subchannel allocation (Figure 9(b)).

Application-level performance. To understand CellFi’s impact on real applications, we model dynamic traffic conditions based on our web workload, and examine web-page download times.

Figure 9 (c) presents the corresponding CDFs of page completion time. The figure highlights that CellFi reduces completion times by 2.3 times at the median compared to Wi-Fi, and roughly by 8% relative to LTE. LTE provides marginally better times at smaller percentiles, however tail performance is significantly degraded due to interference. We also examined whether the network converges.

We observe that the vast majority of access points only hop very few times in all of our runs; roughly 1%-2% of access points do not converge due to interference and hop almost continuously. We omit these figures due to space limitations.

Overheads of signaling. CellFi uses mode 3-0 higher layer configured sub-band CQI feedback reports, which consists of 1 wideband CQI value (4 bits) and 13 sub-band CQI values (2 bits). The payload size for a single mode 3-0 report on a 5 MHz channel is 20 bits per report. The overhead of signaling is 10 Kbps on the uplink for a reporting period of 2 ms.

7 DISCUSSION

In this section, we mention a number of important points related to CellFi design that have not been discussed.

Coexistence between CellFi and 802.11af: CellFi focuses on coexistence among LTE nodes. There are several other efforts (LTE-U, LAA, LWA) that look into coexistence between LTE and WiFi. These are orthogonal solutions that could be deployed along CellFi to enable coexistence with 802.11af.

Centralized vs distributed control plane: CellFi deploys a decentralized control plane. We show that it is efficient and comparable with the state-of-art centralized control plane [20]. We also note that CellFi can be extended to include centralized coordination among nodes from one provider, and distributed coordination across multiple providers, which could further improve performance.

Mobility and roaming: CellFi inherits the benefits of the LTE architecture. It provides seamless roaming across access points, which is difficult to engineer in current WiFi deployments.

Channel aggregation and power optimization: CellFi currently only uses a single TV channel for its operations. One can think of a more flexible channel allocation that will allow channel aggregation and optimization for power. However, these raises other challenges, such as how to detect interference in partially overlapping channels [32], which we leave as future work.

Ease of deployability: CellFi works with unmodified LTE base-band chipsets. However, it does require changes on the access point side. We note that the majority of LTE small cells today are built in software, on programmable reference platforms from TI, Freescale, Broadcom and Qualcomm. Hence we speculate that CellFi can be implemented entirely in software on top of these commodity platforms.

8 RELATED WORK

TV white spaces have recently become available for license exempt use [1, 33]. There are several currently proposed candidate technologies, such as 802.11af [2] and 802.22 [34], but 802.11af seems to be the only one under active development.

LTE extensions have recently been proposed that seek to exploit unlicensed spectrum [35–37], but all these require an anchor, licensed spectrum for the LTE network. The main proposed standards are LAA, LWA and LTE-U. LTE-U focuses only on coexistence between LTE and WiFi using adaptive on/off duty cycling such that effectively only LTE nodes are operational during On period. But it has no mechanism of sharing spectrum between two interfering LTE networks. This is an orthogonal solution that could be deployed along CellFi to enable coexistence with 802.11af. LAA has proposed a contention protocol similar to WiFi CSMA to medium access. This enables LAA networks to coexist with both WiFi and other LAA networks, however its performance will suffer in long range whitespace networks as it will face similar MAC inefficiencies as 802.11af since the medium access mechanism is similar in both technologies. MulteFire [38] is proposed LTE small cell technology for standalone operation in unlicensed spectrum. Similar to LAA, it uses LBT mechanism to do over the air contention as it has been designed keeping in mind the 5 GHz unlicensed spectrum band with LBT requirement in many markets. Therefore in long range whitespace networks it will face similar MAC inefficiencies as 802.11 af.

There are also numerous proposed solutions for coordinating interfering LTE access points, such as SON, ICIC, eICIC [19], but these are vague on protocol details. FERMI [20] proposes a centralized resource management solution in OFDMA networks, however it assumes cooperation among operators which is not realistic in our setting. RADION [39] is a distributed resource management system designed for femtocell networks and does not scale well to large deployments.

WiFi channel allocation has been extensively researched. There have been several studies [40] [41] that address resource allocation using graph coloring for WiFi networks, but their aim is to get maximum number of orthogonal contiguous channels to each interfering AP. Our work aims at getting fair allocation for interfering clients by utilizing as many spectrum fragments (not necessarily contiguous) as possible. Wi-Fi 802.11af can be made to use more than one channel [32], and LTE can achieve this with carrier aggregation. We leave exploring these options for future work.

Several papers proposed the idea to use LTE in TV white spaces [42, 43] but none has described an architecture or proposed an efficient distributed interference management system compatible with today's hardware.

9 CONCLUSIONS

We have designed CellFi, a long-range LTE-based network operating in TVWS. CellFi interfaces with a TVWS spectrum database and uses a fully decentralized interference management algorithm that allows for uncoordinated and unplanned deployment but is compatible with the existing LTE network stack. We show the feasibility of CellFi using experiments with off-the-shelf LTE equipment.

Our extensive simulation results show that CellFi improves MAC performance compared to Wi-Fi and LTE by reducing the number of starved clients by 70%-90%, without affecting the total throughput of the network or penalizing application performance.

REFERENCES

- [1] ETSI EN 301 598. White space devices (WSD); wireless access systems operating in the 470 MHz to 790 MHz TV broadcast band; harmonized EN covering the essential requirements of article 3.2 of the R&TTE directive, 2014.
- [2] 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.
- [3] A. Mody and G. Chouinard. Ieee 802.22 wireless regional area networks. Technical report, 2010.
- [4] Weightless sig. <http://www.weightless.org>.
- [5] Wikipedia. Broadband universal service. https://en.wikipedia.org/wiki/Broadband_universal_service.
- [6] LG Electronics Alcatel-Lucent, Ericsson and etc. Qualcomm. LTE-U SDL coexistence specifications. LTE-U Forum, October 2015.
- [7] Yubing Jian, Chao-Fang Shih, Bhuvana Krishnaswamy, and Raghupathy Sivakumar. Coexistence of wi-fi and laa-lte: Experimental evaluation, analysis and insights. In *Communication Workshop (ICCW), 2015 IEEE International Conference on*. IEEE, 2015.
- [8] P. Surampudi and S. Mohanty. LTE-advanced in white space – a complementary technology. Technical report, 2011. <http://go.radisys.com/rs/radisys/images/paper-lte-white-space.pdf>.
- [9] ETSI TS 136 101. LTE; evolved universal terrestrial radio access (E-UTRA); user equipment (UE) radio transmission and reception (3GPP TS 36.101 version 12.12.0 release 12), 2016.
- [10] Broadcast incentive auction. <https://www.fcc.gov/about-fcc/fcc-initiatives/incentive-auctions>.
- [11] ETSI TS 136 104. LTE; evolved universal terrestrial radio access (E-UTRA); base station (bs) radio transmission and reception (3GPP TS 36.104 version 12.12.0 release 12), 2016.
- [12] ETSI TS 36 423. X2 application protocol (x2ap), release 10), 2013.
- [13] Martin Heusse, Franck Rousseau, Romaric Guillier, and Andrzej Duda. Idle sense: An optimal access method for high throughput and fairness in rate diverse wireless lans. In *Proceedings of the 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM '05*, pages 121–132, New York, NY, USA, 2005. ACM.
- [14] White space database administration. <https://www.fcc.gov/general/white-space-database-administration>.
- [15] White space database operators. <https://tvws-databases.ofcom.org.uk/>.
- [16] IETF. Protocol to access WS database (PAWS), 2014. <https://www.ietf.org/wg/concluded/paws.html>.
- [17] S. Sesia, I. Toufik, and M. Baker. *LTE: The UMTS Long Term Evolution, From Theory to Practice*. Wiley, 2011.
- [18] C. Johnson. *LTE in bullets*. Chris Johnson, 2012.
- [19] H. Holma, A. Toskala, and J. Reunanen. *LTE Small Cell Optimization: 3GPP Evolution to Release 13*. Wiley, 2016.
- [20] 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.
- [21] A.O. Xiaobin Yang, Fapojuwo. Enhanced preamble detection for PRACH in LTE. In *Wireless Communications and Networking Conference (WCNC)*, 2013.
- [22] ETSI TS 136 213. LTE; evolved universal terrestrial radio access (E-UTRA); physical layer procedures (3GPP TS 36.213 version 8.3.0 release 8), 2008.
- [23] Apurv Bhartia, Deeparnab Chakrabarty, Krishna Chintalapudi, Lili Qiu, Bozidar Radunovic, and Ramachandran Ramjee. Iq-hopping: Distributed oblivious channel selection for wireless networks. In *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing, MobiHoc '16*, pages 81–90, New York, NY, USA, 2016. ACM.
- [24] Enterprise and public access. <http://www.ipaccess.com/en/public-access>.
- [25] Nominet. Tv white space database. <http://www.nominet.uk/emerging-technology/tv-white-space/>.
- [26] Broadband Forum. Tr-196: Femto access point service data model. <https://www.broadband-forum.org/technical/download/TR-196.pdf>.
- [27] ns-3 simulator. <https://www.nsnam.org/>.
- [28] J. J. Lee and M. Gupta. A new traffic model for current user web browsing behavior. Technical Report MSR-TR-2014-132, September 2007.
- [29] 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.

- [30] ETSI TS 136 211. LTE; evolved universal terrestrial radio access (E-UTRA); physical channels and modulation (3GPP TS 36.211 version 8.3.0 release 8), 2008.
- [31] How many people does one cell tower usually cover? <https://www.quora.com/How-many-people-does-one-cell-tower-usually-cover-And-how-many-3G-HSDPA-simultaneous-users>.
- [32] 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.
- [33] Second memorandum opinion and order (fcc 10-174). https://apps.fcc.gov/edocs_public/attachmatch/FCC-10-174A1.pdf.
- [34] 802.22. https://en.wikipedia.org/wiki/IEEE_802.22.
- [35] Qualcomm wants LTE deployed in unlicensed spectrum. <http://www.fiercewireless.com/story/qualcomm-wants-lte-deployed-unlicensed-spectrum/2013-11-21>.
- [36] The unlicensed spectrum usage for future IMT technologies – efficient LTE technology enables better performance and experience. Huawei.
- [37] Ericsson first to give smartphone users indoor boost with license assisted access. <http://www.ericsson.com/news/1884277>.
- [38] Multefire release 1.0 technical paper: A new way to wireless. <https://www.multefire.org/wp-content/uploads/MulteFire-Release-1.0-whitepaperFINAL.pdf>.
- [39] 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.
- [40] 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.
- [41] 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.
- [42] Mihaela C. Beluri, Erdem Bala, Yuying Dai, Rocco Di Girolamo, Martino Freda, Jean-Louis Gauvreau, Scott Laughlin, Debashish Purkayastha, and Athmane Touag. Mechanisms for lte coexistence in tv white space. In *DySPAN*, pages 317–326. IEEE, 2012.
- [43] P. Surampudi and S. Mohanty. Lte-advanced in white space – a complementary technology. <http://go.radisys.com/rs/radisys/images/paper-lte-white-space.pdf>.