Before the
Federal Communications Commission
Washington, D.C. 20554

In the Matter of	)	
	)	
Office of Engineering and Technology	)	ET Docket No. 15-105
and Wireless Telecommunications	)	
Bureau Seek Information on Current	)	
Trends in LTE-U and LAA Technology	)	
	)	

COMMENTS OF GOOGLE INC.

In keeping with its restrained regulatory approach to unlicensed use, the Commission has established straightforward technical rules for unlicensed spectrum, while leaving the development of coexistence mechanisms to industry cooperation rather than regulatory intervention. This has worked well, with diverse technologies developing ways to use the same spectrum bands effectively.

Interest in deploying the Long-Term Evolution (LTE) standard over unlicensed spectrum is further evidence of the success of the Commission's light-touch approach, and a further demonstration that unlicensed frequencies support innovation. Yet to the extent that the use of LTE in unlicensed spectrum depends on having access to both unlicensed *and* licensed spectrum, this technology presents novel coexistence challenges. The attached white paper by Nihar Jindal and Donald Breslin, *LTE and Wi-Fi in Unlicensed Spectrum: A Coexistence Study*, summarizes Google's initial investigation into the issue of coexistence between one version of unlicensed LTE technology (LTE-Unlicensed, or LTE-U) and Wi-Fi.[1] The paper shows that in many circumstances, LTE-U coexists poorly with Wi-Fi in the 5 GHz band.

---

[1] *See* Attachment A, Nihar Jindal & Donald Breslin, *LTE and Wi-Fi in Unlicensed Spectrum: A Coexistence Study* (2015) (Jindal & Breslin Study).

Two aspects of LTE-U, in particular, affect Wi-Fi transmissions. First, LTE-U's approach to duty-cycling causes LTE transmissions to begin abruptly, often in the middle of Wi-Fi transmissions, interrupting them and causing Wi-Fi to slow down throughput in response to increased error rates.[2] Second, Wi-Fi operation in the presence of moderate interference from LTE (below Wi-Fi's prescribed energy detect level) experiences substantial degradation.[3]

In order to preserve unlicensed spectrum as a platform for permissionless innovation, LTE operators must work with other users of unlicensed bands to overcome these technical issues and ensure that license-anchored systems will not systematically crowd out popular technologies that rely solely on unlicensed spectrum.

1. A Balance of Licensed and Unlicensed Spectrum Fuels the Wireless Economy.

In the last 20 years, radio technologies that rely on unlicensed spectrum, such as Bluetooth, Radio Frequency Identification (RFID),[4] Wi-Fi, and ZigBee,[5] have flourished. Indeed, access to both unlicensed and licensed spectrum has enabled tremendous innovation in wireless devices and services. As the Commission has recognized, "[b]ecause they are free from the delays inherent in the licensing process, unlicensed devices can frequently be designed to fill a unique need and be introduced into the marketplace rather quickly."[6] A balanced spectrum policy has allowed expansive growth of the wireless economy, benefiting consumers, innovators, and investors. Exclusive

---

[2] *Id.* at Section 3.1.

[3] *Id.* at Section 3.2.

[4] RFID technologies are used in a variety of industries to track inventory or other objects.

[5] Zigbee powers technologies that benefit from ad hoc and mesh networking solutions, such as home automation.

[6] Kenneth R. Carter, Ahmed Lahjouji, & Neal McNeil, FCC, *Unlicensed and Unshackled: A Joint OSP-OET White Paper on Unlicensed Devices and Their Regulatory Issues*, OSP Working Paper Series at 5 (May 2003).

access to licensed spectrum has provided the certainty major operators need to make large, long-term investments in their wide-area networks, while broad eligibility for access to unlicensed spectrum fosters widespread contributions to innovation as well as fast-paced investment in emerging technologies and smaller, more distributed networks.

2. Interest in Deploying LTE over Unlicensed Frequencies Further Shows that Unlicensed Spectrum Is Attractive for a Wide Range of Technologies.

In recent months, Qualcomm, Ericsson, Verizon, T-Mobile, and others have expressed interest in deploying LTE over unlicensed spectrum.[7] This interest further demonstrates that unlicensed bands continue to be attractive for a wide variety of new technologies and applications. Both prominent variations of LTE over unlicensed rely on access to both licensed *and* unlicensed spectrum. The LTE control channels operate in a licensed frequency, and a mobile operator has the flexibility to send data to devices over licensed and/or unlicensed spectrum. Thus, if the unlicensed frequencies becomes congested, a control channel may instruct the device to revert to licensed spectrum for data transmission, or vice-versa. This arrangement provides licensed operators access to additional spectrum without the associated expense of obtaining a license, while still allowing them to maintain the high quality of service expected by consumers for such licensed services.[8] These uses of LTE are fundamentally license-anchored: although they

---

[7] *See, e.g.*, LTE-U Forum, http://www.lteuforum.org/ (last visited June 11, 2015) (listing Alcatel-Lucent, Ericsson, Qualcomm Technologies, Inc., a subsidiary of Qualcomm Incorporated, Samsung, and Verizon as members); Mike Dano, *Verizon, T-Mobile Push Unlicensed LTE Forward - But Concerns Remain*, FierceWireless, Mar. 3, 2015, http://www.fiercewireless.com/story/verizon-t-mobile-push-unlicensed-lte-forward-concerns-remain/2015-03-03.

[8] *See, e.g.*, Qualcomm, *Making the Best Use of Unlicensed Spectrum for 1000x*, at 5, 6 (Feb. 2015) (referring to this use as having a license-anchor and noting that licensed spectrum is “the foundation” on which more efficient use of unlicensed spectrum is built), https://www.qualcomm.com/documents/making-best-use-unlicensed-spectrum-presentation.

operate in unlicensed spectrum, they cannot be used without access to licensed spectrum.

There are currently two well-understood efforts to develop technical specifications for use of LTE in unlicensed spectrum—LTE-U and License-Assisted Access (LAA). LTE-U supports simultaneous coordinated operation across licensed and unlicensed bands, is being developed by an industry consortium (the LTE-U Forum), and targets operation in the 5 GHz unlicensed spectrum.[9] LTE-U is designed for countries without a listen-before-talk (LBT) requirement, such as the United States.[10] LAA is the 3rd Generation Partnership Project (3GPP) effort to standardize operation of LTE in unlicensed bands.[11] Unlike LTE-U, LAA is expected to perform LBT,[12] as required by certainly regulatory regimes, including those in Europe. The European Telecommunications Standards Institute (ETSI) supports the development of standards applicable in European jurisdictions and has established specific requirements for LBT.

3. LTE in Unlicensed Spectrum Presents Challenges in Coexisting with Other Users of the 5 GHz Unlicensed Bands.

Coexistence among diverse technologies in the unlicensed bands has typically been resolved through cooperation and without regulatory intervention. For example, the industry standards for Wi-Fi and Bluetooth incorporate mechanisms to ensure that they can fairly coexist with each other.[13]

---

[9] *See* LTE-Forum, http://www.lteuforum.org/.

[10] Mitchell Shapiro, *Some Clarification on LTE in Unlicensed Spectrum*, Quello Center, Apr. 2, 2015, http://quello.msu.edu/some-clarification-on-lte-in-unlicensed-spectrum/.

[11] Dino Flore, *Evolution of LTE in Release 13*, 3GPP, Feb. 18, 2015, http://www.3gpp.org/news-events/3gpp-news/1628-rel13.

[12] Phil Goldstein, *AT&T in No Hurry to Test and Deploy LTE-Unlicensed*, FierceWirelessTech, Apr. 1, 2015, http://www.fiercewireless.com/tech/story/att-no-hurry-test-and-deploy-lte-unlicensed/2015-04-01; s*ee also* Jindal & Breslin Study, Appendix 2.

[13] *See* Jindal & Breslin Study at 18 n.24.

As currently conceived, however, LTE-U and LAA present new challenges for coexistence with other unlicensed technologies. The accompanying Jindal & Breslin Study summarizes an initial investigation into the issue of coexistence between LTE-U and Wi-Fi in the 5 GHz band. The paper shows that in many circumstances, LTE-U coexists poorly with Wi-Fi. The failure to coexist effectively can be attributed to two factors: (1) the effect of LTE-U's duty-cycling mechanism on Wi-Fi operation and (2) the lack of effective coexistence mechanisms in scenarios where LTE-U and Wi-Fi devices receive each other's signals at moderate levels.[14]

In some cases, LTE-U uses duty-cycling to enable time-sharing between LTE-U and Wi-Fi devices.[15] The premise is that the Wi-Fi link will operate at its interference-free data rate during the LTE-U OFF cycle, and LTE-U will operate at its interference-free rate during the LTE-U ON cycle.[16] One might expect this regime to result in each system getting its interference-free throughput during the time each system is on; for example, if one LTE-U access point shared a channel with one Wi-Fi access point, and the LTE-U access point employed a 50% duty cycle, then equitable sharing would result.[17] However, as the Jindal & Breslin Study shows, LTE-U begins transmitting at the start of every ON cycle without checking if the medium is clear.[18] As a result, the start of an LTE-U transmission commonly interrupts a Wi-Fi transmission in mid-frame and causes that frame to be received in error.[19] As the frame error rate increases, a Wi-Fi device generally responds

---

[14] *Id.* at Sections 3.1, 3.2.

[15] As discussed in further detail below, because the LTE-U Forum defines duty-cycling tests for settings where Wi-Fi is received at a level above -62 dBm, it is not clear whether LTE-U will duty-cycle when Wi-Fi is heard below -62 dBm.

[16] Jindal & Breslin Study at 6.

[17] *See id.*

[18] *Id.*

[19] *Id.* at 6-7.

by reducing its transmission rate and, therefore, its throughput.[20] For instance, when an LTE-U device operates on a 50% duty cycle, Wi-Fi speeds are slowed to one-quarter of their interference-free rate.[21]

Equally importantly, coexistence between Wi-Fi and LTE-U devices in unlicensed spectrum is likely to be poor when those devices receive interfering signals at moderate power levels. Below we first discuss lower-power or distant LTE operation, and then lower-power or distant Wi-Fi operation.

Wi-Fi Devices Receive LTE Signals at Power Levels Below -62 dBm. Wi-Fi devices will not transmit while LTE-U signals are received at a power above -62 dBm because Wi-Fi's energy detect threshold is exceeded. As discussed above, that characteristic limits Wi-Fi's access to the 5 GHz band.[22] Surprisingly, however, signals that are quieter than -62 dBm may cause even more severe interference to Wi-Fi. If the signal received by a Wi-Fi device falls below the -62 dBm threshold, then the Wi-Fi devices will not wait to transmit, even in the presence of significant interference not formally "detected" by the Wi-Fi device. Transmitting during LTE-U's ON time, rather than only during its OFF time, leads to greater error rates and causes Wi-Fi to slow down even more than in the tests described above.[23]

LTE-U Devices Receive Wi-Fi Devices at Power Levels Below -62 dBm. Regardless of the power at which an LTE-U device is transmitting, coexistence problems may arise if the LTE-U device fails to use a duty cycle when it hears Wi-Fi at a signal level below -62

---

[20] *Id.* at 7.

[21] *Id.* at 9.

[22] *Id.* at 6.

[23] *Id.* at 15-16.

dBm.[24] It is unclear whether an LTE-U device will use duty-cycling at all in these conditions because the LTE-U Forum only defines duty-cycling tests for settings above -62 dBm.[25] The LTE-U device may transmit continuously. In that situation, Wi-Fi transmissions are likely to be substantially impaired by the continuous LTE transmission and, in some situations, Wi-Fi will fail completely.

Many of the coexistence challenges presented by LTE-U apply equally to LAA. As noted above, LAA is expected to make modifications to LTE to meet ETSI's listen-before-talk requirement.[26] While LBT is a key element of Wi-Fi's coexistence protocol, it is not, by itself, a guarantee of effective sharing with other users of the 5 GHz band. As described in the Jindal & Breslin Study, additional factors will also affect coexistence. Two factors are worth highlighting specifically.

First, the specific responses of LTE and Wi-Fi devices after an unlicensed channel is detected as "clear" will dramatically affect coexistence. Wi-Fi incorporates several distinct coexistence mechanisms. Initially, Wi-Fi uses LBT: it listens and waits for the channel to be clear.[27] Once the channel is clear, Wi-Fi waits for an additional, randomly selected interval, and then, only if the channel is still clear does it begin transmission.[28] This is

---

[24] The LTE-U Forum defines duty-cycling tests exclusively for above -62 dBm settings. LTE-U Forum, *LTE-U SDL Coexistence Specifications* V1.0 § 5.1 (2015-02), http://www.lteuforum.org/uploads/3/5/6/8/3568127/lte-u_forum_lte-u_sdl_coexistence_specifications_v1.0.pdf.

[25] *Id.*

[26] Flore, *supra* note 10; *see also* Goldstein, *supra* note 11. Listen-before-talk requires a device to wait until the measured receive power is below a defined threshold (-60 dBm per the relevant ETSI standards for 20 MHz operation) before beginning transmission. ETSI, *Broadband Radio Access Networks (BRAN); 5 GHz High Performance RLAN; Harmonized EN Covering the Essential Requirements of Article 3.2 of the R&TTE Directive*, ETSI EN 301 893 v1.8.1, § 4.8 (Mar. 2015).

[27] Jindal & Breslin Study at 24.

[28] *Id*.

known as “random backoff.” If a collision—e.g., simultaneous transmissions by Wi-Fi devices—occurs, then larger random intervals are chosen before the next transmission to decrease the probability of collision.[29] This increase in length of the random intervals is referred to as "exponential backoff."[30] Exponential backoff is critical for coexistence among Wi-Fi devices because it decreases the likelihood of collisions even in a crowded medium, and as a result, the precise method for implementing it could have substantial implications on effective operation of Wi-Fi in the presence of LAA.[31] The interplay of LBT, random backoff, and exponential backoff determine the degree to which Wi-Fi makes room for other users of unlicensed bands, and its vulnerability to interference from those other users.

Second, the ETSI standard to which LAA is being tailored does not require devices to perform LBT if they receive signals at moderate levels. Rather, the standard only requires LBT be performed to a threshold of -60 dBm. The Jindal & Breslin Study, though, demonstrates that interference to Wi-Fi signals weaker than -60 dBm can still be extremely harmful.[32] This limitation could effectively permit LAA devices to transmit continuously so long as Wi-Fi or other signals are received at a level below -60 dBm.[33] Operation of LTE in such conditions can degrade a Wi-Fi link—including by causing a complete loss of Wi-Fi connectivity.[34]

---

[29] *Id.*

[30] *Id.*

[31] *Id.* at 17.

[32] *Id.* at Section 3.2.

[33] *Id.* at 17.

[34] *Id.*

4. Developing Effective Coexistence Mechanisms Between License-Anchored Technologies and Unlicensed-Only Technologies Is Critical for Preserving Unlicensed Spectrum as a Platform for Innovation.

Coexistence between unlicensed technologies is critically important to maximizing opportunities for beneficial use of this shared spectrum. The coexistence challenges presented by both LTE-U and LAA in the 5 GHz band are especially difficult, because these license-anchored systems can shift traffic from unlicensed to licensed spectrum whenever interference on unlicensed spectrum becomes unacceptable. Put differently, LTE-U and LAA operators will have an exclusive, licensed spectrum option that allows them to minimize or entirely avoid the ill-effects of poor coexistence in the 5 GHz band, while other users of the same unlicensed band must rely on cooperative solutions.

When multiple unlicensed-only technologies operate in a band, respectful coexistence is in the best interest of each user, because the alternative interference is generally mutually destructive. By contrast, because a mobile operator deploying LTE-U or LAA in unlicensed spectrum does not face the risk of fatal interference to its service, it may have less incentive to avoid conflict with purely unlicensed users. The problem is particularly worrisome inasmuch as licensed mobile operators may view some Wi-Fi providers, such as cable companies offering Wi-Fi hotspots to their customers, as competitors.

If, because of their divergent incentives, operators of LTE-U and LAA fail to develop means for fair coexistence with other technologies, license-anchored technologies could crowd out unlicensed-only systems. Ultimately, access to licensed frequencies could become a de facto prerequisite for successful use of the unlicensed spectrum where license-anchored services operate.

Allowing that to happen would undermine the benefits of enabling unlicensed use. Today, unlicensed users can affordably and nimbly deploy new technologies and products without seeking permission from license holders or a regulator, so long as they adhere to technical rules. Effectively requiring access to licensed spectrum would increase the cost of developing new unlicensed technologies and the time required to bring them to market. Indeed, the licensed and unlicensed models could be collapsed into one model that has the economic characteristics of today's licensed model of deployment, undermining the balanced spectrum that has enabled the wireless revolution of the last 20 years.

Conclusion

In order to maintain the balanced spectrum policy that has benefited innovators, entrepreneurs, businesses, and consumers, the Commission should be vigilant in ensuring that deployments of LTE-U and LAA in unlicensed spectrum will not systematically exclude unlicensed-only technologies.

Respectfully submitted,

Aparna Sridhar

Austin C. Schlick
Director, Communications Law
Aparna Sridhar
Counsel
Google Inc.
25 Massachusetts Avenue NW
Ninth Floor
Washington, DC 20001

June 11, 2015

ATTACHMENT A

LTE and Wi-Fi in Unlicensed Spectrum: A Coexistence Study

Nihar Jindal & Donald Breslin

Table of Contents

Abstract

This work investigates the coexistence between LTE technology operating simultaneously across licensed and unlicensed spectrum and Wi-Fi in the 5 GHz unlicensed bands. Two versions of such LTE technology are considered: LTE-U, which is a proprietary evolution of LTE intended for regions without a listen-before-talk (LBT) requirement, and License Assisted Access (LAA), which is 3GPP's ongoing effort to standardize simultaneous operation of LTE in licensed and unlicensed bands. LTE-U is shown to coexist poorly with Wi-Fi primarily due to two factors: (1) the incompatibility of LTE-U's duty-cycling mechanism with Wi-Fi equipment and (2) the lack of an effective coexistence mechanism in scenarios where LTE-U and Wi-Fi devices hear each other at moderate but non-negligible power levels. LAA will perform LBT as required by certain national and regional regulations. Our investigation reveals, however, that inclusion of LBT does not by itself guarantee successful coexistence with Wi-Fi and other purely unlicensed technologies.

Underlying the coexistence challenge is the fact that LTE-U and LAA are license-anchored systems that can fluidly move from unlicensed to licensed bands whenever interference on the unlicensed bands becomes unacceptable. This may reduce the incentive of these systems to coexist in the unlicensed bands.

Section 1: Introduction

Access to unlicensed spectrum has been the key to the development and widespread adoption of innovative wireless technologies such as Wi-Fi, Bluetooth, and ZigBee. Indeed, Cisco estimates that nearly half of all worldwide Internet traffic is carried on unlicensed bands using Wi-Fi technology, compared to only 4% on mobile wide-area networks.[1] LTE, which is the leading mobile wireless technology, is broadly deployed, although exclusively in licensed bands to date. *LTE in unlicensed* is an evolution of LTE technology that enables simultaneous operation across licensed and unlicensed bands, and thus LTE is emerging as a new and potentially powerful technology for the unlicensed bands. LTE-U is the version of this technology developed by an industry consortium (the LTE-U Forum), while License Assisted Access (LAA) is 3GPP's ongoing effort to standardize simultaneous operation across licensed and unlicensed bands as part of LTE Release 13. We use the term "LTE in unlicensed" to refer to both LTE-U and LAA collectively.

LTE in unlicensed is a *license-anchored* system, meaning that it is operated by holders of exclusively licensed spectrum on licensed and unlicensed bands simultaneously: licensed spectrum is used—due to its guaranteed availability—for transmission of control and QoS traffic, while unlicensed spectrum is used for best-effort data. According to developers of this technology, this approach provides mobile operators with higher spectral efficiency and increased control, and streamlines management compared to Wi-Fi offloading. In these respects, LTE in unlicensed provides another example of the innovation unlocked by unlicensed spectrum.

Given how ubiquitous use of unlicensed bands has become, a critical question for any new technology entering these bands is how well it coexists with other deployed technologies. In the first part of this paper (Sections 2-4), we report on an investigation of the impact the two variants of LTE in unlicensed (LTE-U and LAA) have on Wi-Fi, which is the most widely deployed technology in the 5 GHz unlicensed bands.

LTE-U targets deployment prior to 3GPP standardization of LTE in unlicensed in countries without listen-before-talk (LBT) requirements, such as the United States.[2] Our experimental investigation revealed the following two principal limitations of LTE-U/Wi-Fi coexistence:

1. LTE-U's primary coexistence mechanism—duty-cycling (i.e., cycling LTE-U through on/off periods) when a Wi-Fi access point is detected at a level above -62 dBm—can reduce the throughput of Wi-Fi devices significantly below the reduction associated with fair sharing, as well as significantly increase Wi-Fi latency.

---

[1] Cisco, *VNI Forecast Highlights* 2015, http://www.cisco.com/web/solutions/sp/vni/vni_forecast_highlights/index.html (last visited June 11, 2015) (figures appear "Wired Wi-Fi and Mobile Growth" section).

[2] *See, e.g.*, LTE-U Forum, http://www.lteuforum.org/ (last visited June 11, 2015) (listing Alcatel-Lucent, Ericsson, Qualcomm Technologies, Inc., a subsidiary of Qualcomm Incorporated, Samsung, and Verizon as members).

2. LTE-U does not have an effective coexistence technique to handle scenarios in which LTE-U and Wi-Fi devices hear each other at moderate (below -62 dBm) power levels and, as a consequence, Wi-Fi can be crippled in such scenarios.

The 3GPP standardization body is actively working to standardize LTE operation in unlicensed bands as part of the overall LTE release 13 effort, and in particular is studying an LAA standard that would meet the LBT requirement of the European regulatory bodies (based on standards derived by the European Telecommunications Standards Institute, or ETSI[3]).[4] LBT is the principle behind Wi-Fi's medium access control, and is employed as well in other unlicensed technologies such as ZigBee, so the inclusion of LBT is a positive step from the coexistence perspective. But LBT is not by itself a guarantee of effective sharing with other users of the 5 GHz bands, and considerable further study is required. Of particular importance are the specifics of how devices perform an additional random backoff after the power on the medium falls below the LBT threshold, which is a critical factor in controlling access to the medium for every device. An additional concern is that ETSI specifications only require LBT be performed to a threshold of -60 dBm, meaning that LAA shares with LTE-U the previously stated concern regarding the lack of coexistence mechanisms for scenarios in which LTE-U and Wi-Fi devices hear each other at a level below -62 dBm.

The second part of this paper (Section 5) views coexistence from a policy perspective, and in particular considers the incentives that drive coexistence. When multiple unlicensed-only technologies operate in a band, respectful coexistence is in the best interest of each operator, because the absence of suitable coexistence procedures is generally *mutually* self-destructive. The license-anchored operation of LTE in unlicensed, however, fundamentally changes the coexistence paradigm: because a license-anchored system operates simultaneously in licensed and unlicensed bands and can dynamically move traffic between the bands on a granular basis (e.g., per-user and per-flow), such a system is inherently less sensitive to collisions and congestion in the unlicensed bands than is a system operating solely in unlicensed spectrum. This may reduce the incentive for a license-anchored system to develop effective coexistence mechanisms in the unlicensed band.

Section 2: Background

This section provides a high-level overview of LTE in unlicensed and Wi-Fi, with more detailed material deferred to the appendices. LTE in unlicensed uses LTE's carrier aggregation (LTE-CA) feature, which allows for operation across multiple (possibly discontiguous) bands, to operate simultaneously in licensed and unlicensed bands. One of these bands acts as the anchor and is designated as the Primary Serving Cell (PCell), which is used to communicate security and control information, while all other bands are referred to as Secondary Serving

---

[3] ETSI is not a regulatory organization, but it does support regulation and legislation (in Europe especially) with technical standards and specifications. *See* www.etsi.org/technologies-clusters/technologies/regulation-legislation.

[4] LBT requires a device to wait until the measured receive power is below a defined threshold (-60 dBm per ETSI standards for 20 MHz operation), before beginning transmission. The standards are documented in ETSI EN 301 893 v1.8.1.

Cells (SCells). LTE-CA is currently defined by 3GPP solely for LTE bands in licensed spectrum, and is already in deployment. LTE in unlicensed extends LTE-CA to unlicensed spectrum: the PCell must remain in a licensed LTE band, while the SCell(s) can be in the 5 GHz unlicensed bands. Control and high-QoS traffic is expected to be carried on the PCell, while the SCell(s) are primarily used for best-effort data. A key property is that traffic can be dynamically moved, on a per-user and per-flow basis, between and across bands. This can be done as dictated by instantaneous interference conditions. For example, an LTE client suffering from high interference on the unlicensed SCell(s) could be served entirely on licensed bands, while other client devices continue to be served on licensed and unlicensed. LTE-U and LAA are the two variants of LTE in unlicensed, and are detailed below. Note that LTE-U and LAA are both restricted to entities with LTE networks deployed in licensed spectrum, and thus LTE-U and LAA do not allow standalone LTE in unlicensed spectrum.

LTE-U refers specifically to the proprietary technology developed by the LTE-U Forum industry consortium. LTE-U is being developed outside of 3GPP and prior to 3GPP's LTE in unlicensed effort (LAA). LTE-U is intended for regulatory regimes without an LBT requirement, such as the United States, and is intended to work without modification to Rel. 10/11/12 LTE. Two major U.S. carriers have made public announcements of plans to deploy LTE-U in late 2015/early 2016.[5] LTE-U defines the PCell to be in a licensed band, and allows for one or two SCells, each 20 MHz, in the 5 GHz unlicensed band (specifically, within the U-NII-1 and/or U-NII-3 bands, which span 5150-5250 MHz and 5725-5825 MHz, respectively). The PCell serves bi-directional traffic but the SCells are downlink-only, meaning that an LTE-U eNodeB (LTE's term for a base station; abbreviated as eNB) transmits in licensed and unlicensed bands, but LTE-U client devices transmit only in licensed bands.

The LTE-U Forum defines three coexistence mechanisms: channel selection, duty-cycling, and opportunistic cell switch-off. Channel selection refers to an LTE-U eNB attempting to choose a channel—for the SCell(s)—on which no Wi-Fi access points (Wi-Fi's term for a base station; abbreviated as AP) or other LTE-U eNBs are detected. When a clean channel cannot be found and LTE-U must operate its SCell(s) on the same channel as Wi-Fi APs and/or other LTE-U eNBs, then the LTE-U duty-cycles its transmissions on the SCells, i.e., cycles through transmit on/transmit off periods. (The PCell in the licensed band is always on.) Finally, the SCell is to be turned completely off whenever LTE traffic demands can be met on the PCell. These mechanisms are specified at a high-level, with details left to implementers. Further detail on LTE-U is provided in Appendix 1, including a list of defined coexistence tests.

LAA is the yet-to-be-finalized 3GPP standardized version of LTE in unlicensed, with completion expected in the first half of 2016. LAA is quite similar to LTE-U, with the key difference that LAA is being designed to meet ETSI's LBT requirements. LBT requires each device to defer transmission until the on-air received energy falls below a specified threshold (-60 dBm for 20 MHz operation, per ETSI requirements); in other words, a device must wait for

---

[5] Mike Dano, *Verizon's LTE-U Forum Completes LTE Unlicensed Tests on Way Toward Carrier's 2016 Deployment in 5 GHz, 3.5 GHz,* FIERCEWIRELESSTECH, Mar. 3, 2015, http://www.fiercewireless.com/tech/story/verizons-lte-u-forum-completes-lte-unlicensed-tests-way-toward-carriers-201/2015-03-03.

any nearby transmissions to complete before beginning its own transmission.[6] An LAA eNB will therefore operate in a bursty and opportunistic manner, transmitting whenever other nearby devices are not transmitting, as opposed to with the fixed duty-cycle approach of LTE-U. Because LTE was designed for licensed spectrum and a centralized management (i.e., network-controlled) model, it is generally an "always-on" technology; as a result, adapting to LBT is a marked change for the LTE protocol. An additional difference is that LAA may allow bi-directional traffic on SCells in unlicensed bands, whereas LTE-U is downlink-only in unlicensed bands.

Unlike LTE, Wi-Fi was originally conceived for use in unlicensed spectrum in which there is no explicit coordination among devices. As a result, the underlying protocol has built into it decentralized spectrum sharing mechanisms. For this reason, the Wi-Fi protocol is based on LBT, although with a few key additional features that go beyond ETSI's LBT requirements.[7] First, Wi-Fi's version of LBT defers to signals that are much weaker than the minimum level required by ETSI. ETSI LBT requires a transmitter to defer if the received energy is above -60 dBm (for 20 MHz), while Wi-Fi defers if the received energy is above -62 dBm (this level is referred to as the energy detect threshold, or ED for short) *or* if a valid Wi-Fi preamble is detected. Wi-Fi's ED threshold is nearly the same as ETSI's LBT threshold, but Wi-Fi preamble detection is required to work to at least -82 dBm, and in reality works to -90 dBm or lower in most products. Hence, Wi-Fi devices defer to other Wi-Fi transmissions much more conservatively (i.e., at a much larger distance) than a device which only meets ETSI requirements.[8] Second, Wi-Fi goes beyond the ETSI requirements in specifying how long a device must wait after the on-air energy falls below the threshold before initiating a transmission. These seemingly subtle differences turn out to be substantial in terms of LAA/Wi-Fi coexistence, as explained in Section 4.

Section 3: Evaluation of LTE-U/Wi-Fi Coexistence

In order to measure the effectiveness of LTE-U's coexistence mechanisms, a series of tests was conducted of retail Wi-Fi equipment operating in the presence of emulated LTE-U transmissions. Since LTE-U equipment is not yet available, LTE-U transmissions were emulated with a signal generator, using the description of LTE-U coexistence mechanisms contained in LTE-U Forum documentation. Because Wi-Fi operates exclusively in unlicensed spectrum and LTE-U is downlink-only, only LTE-U transmissions in the unlicensed bands are considered and LTE-U client devices need not be emulated. All tests were conducted over-the-air, in an RF isolation chamber with programmable attenuators that allow the Wi-Fi/Wi-Fi and Wi-Fi/LTE-U link qualities to be controlled. In order to understand the fundamental nature of Wi-Fi/LTE-U interactions, a single Wi-Fi AP-client pair running a TCP or UDP session (using iperf, an industry standard throughput testing tool) in the face of interference from a single LTE-U eNB, with all devices operating on a single 20 MHz channel in

---

[6] *See* Appendix 2 for a description of ETSI's LBT regulations.

[7] *See* Appendix 3 for a description of the Wi-Fi protocol.

[8] The 22 dB difference in terms of power threshold (-82 dBm for Wi-Fi preamble detection versus -60 dBm per ETSI) corresponds to more than a factor of 10 in terms of distance in free space.

the U-NII-3 band, was studied. Wi-Fi devices that support one, two, and three spatial streams were tested. Testing focused on the effect of LTE-U duty-cycling upon Wi-Fi when all devices are in close range (above the -62 dBm ED threshold) to each other, and the effect of LTE-U when the LTE-U eNB and Wi-Fi AP hear each other at a power below -62 dBm. Fundamental limitations were discovered in both regimes, as detailed in the following subsections.

Section 3.1: LTE-U Duty-Cycling Is Incompatible with Wi-Fi Devices in 5 GHz

The principal mechanism by which LTE-U interacts with co-channel Wi-Fi devices (received by the LTE-U eNB at a power above -62 dBm) is to duty-cycle the LTE-U SCell(s) operating in the 5 GHz unlicensed band: the LTE-U eNB completely shuts off all SCell transmissions[9] for X milliseconds (ms), followed by a burst of continuous LTE transmissions for Y ms, and so forth. X and Y refer to the LTE-U OFF and ON time, respectively, while the quantity X+Y is the total duty cycle period. The duty cycle refers to the percentage of time that LTE-U is on, and thus is Y/(X+Y).

Because LTE-U duty-cycling coexistence tests (Appendix 1) are defined only for scenarios in which the LTE-U eNB receives the Wi-Fi AP(s) at a power level above -62 dBm, we focus only on such a scenario in this subsection. In the case of a single co-channel Wi-Fi AP with full-buffer traffic, the duty cycle is constrained by the LTE-U Forum's coexistence tests to be no larger than 50%, and the LTE-U ON time Y is constrained to be no larger than 50 ms; if sharing with 2 co-channel full-buffer APs, then the duty-cycle constraint becomes 33%. No tests are defined with more than 2 co-channel Wi-Fi APs.

Duty-cycling purports to enable efficient time-sharing between LTE-U and Wi-Fi devices, under the premise that the Wi-Fi link will operate at its interference-free data rate during the LTE-U OFF cycle, and LTE-U will operate at its interference-free rate during the LTE-U ON cycle. Wi-Fi devices will not transmit while LTE-U is on because Wi-Fi's ED threshold (-62 dBm) is exceeded. Based on this line of thinking, one would expect each system to get its interference-free throughput scaled by the fraction of time each system is on, i.e., Wi-Fi and LTE-U throughput equal to $R_{Wi\text{-}Fi}$ * X/(X+Y) and R_{LTE}* Y/(X+Y), respectively, where $R_{Wi\text{-}Fi}$ and R_{LTE} correspond to the Wi-Fi/LTE interference-free throughputs. Due to the fact that the LTE-U device explicitly controls the ON/OFF cycle, this throughput expectation is by design met for LTE-U. The situation is quite different for Wi-Fi, however, because Wi-Fi devices are not explicitly aware of the duty-cycled nature of the LTE-U interferer nor is the Wi-Fi protocol designed for such interference.

Our results show that LTE-U's duty-cycling does not necessarily lead to efficient time-sharing and, as a result, Wi-Fi often achieves a throughput significantly lower than the above predictions. As explained in detail below, the fundamental reason for degraded Wi-Fi throughput is the fact that LTE-U begins transmitting at the start of every ON cycle without checking if the medium is clear. As a result, the start of every LTE-U transmission almost

---

[9] The LTE-U Forum makes an exception for Rel. 12 Discovery signals, which can be 1-5 ms in duration, and can be repeated as often as every 40 ms. These signals are allowed to be transmitted during the nominal LTE-U OFF time, but are not counted towards the ON duration.

inevitably interrupts a Wi-Fi transmission in mid-frame and causes that frame to be received in error; such an error occurs every duty-cycle period, and often causes the transmitting Wi-Fi device to reduce its rate (and thus the Wi-Fi link throughput).

Shown below is a capture of the received power (at a signal analyzer) versus time, for a Wi-Fi (using IEEE 802.11ac) UDP session operating in the presence of an LTE-U eNB using a 30 ms ON/30 ms OFF duty-cycle. All devices hear each other well above the -62 dBm threshold. The UDP session is in the uplink (client to AP) direction, so the only downlink frames are short Wi-Fi acknowledgements (ACK's) sent by the AP. To help parse the plot, note that the signal analyzer is positioned so that transmissions from the Wi-Fi AP, Wi-Fi client, and LTE-U eNB correspond to received power of (approximately) 0, -5, and -20 dBm (y-axis). Note also that "int", short for interference, refers to LTE-U transmissions in the three figures below.

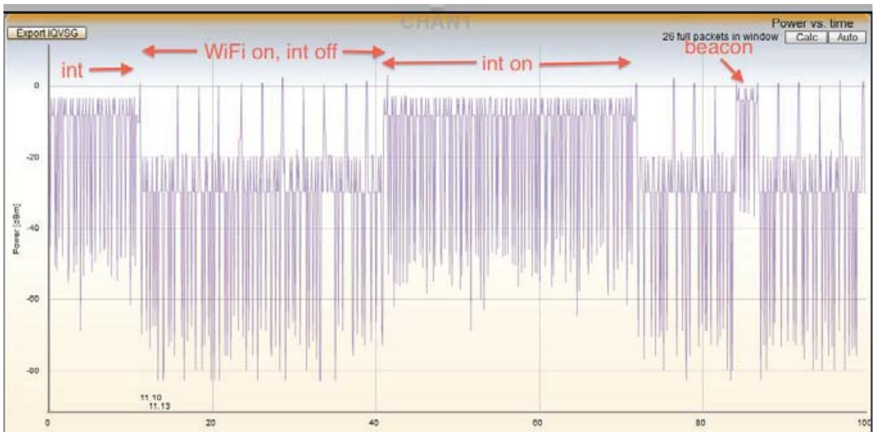

At the beginning of the capture, LTE-U transmissions (but no Wi-Fi) are seen for the first approximately 10 ms. Once LTE-U transmission stops, Wi-Fi detects the channel as clear and begins transmission. A number of long uplink frames are seen, interspersed by short downlink ACK's (seen as small vertical spikes going up to 0 on the y-axis). LTE-U begins transmitting again shortly after 40 ms, Wi-Fi turns off in a manner detailed below, and then Wi-Fi starts again when LTE-U ends transmission. This figure confirms that Wi-Fi and LTE-U alternate transmission, as expected when duty-cycling is performed and LTE-U received by the Wi-Fi devices at a power level above Wi-Fi's ED threshold.

The next figure zooms in on the end of an LTE-U ON cycle. As expected, once LTE-U ceases transmissions, Wi-Fi detects the medium as clear and performs a backoff (of approximately 150 microseconds in this particular case), after which Wi-Fi begins transmitting frames. In this particular case, the transmissions begin with an RTS/CTS (request-to-send and clear-to-send) exchange following by a long uplink data frame.

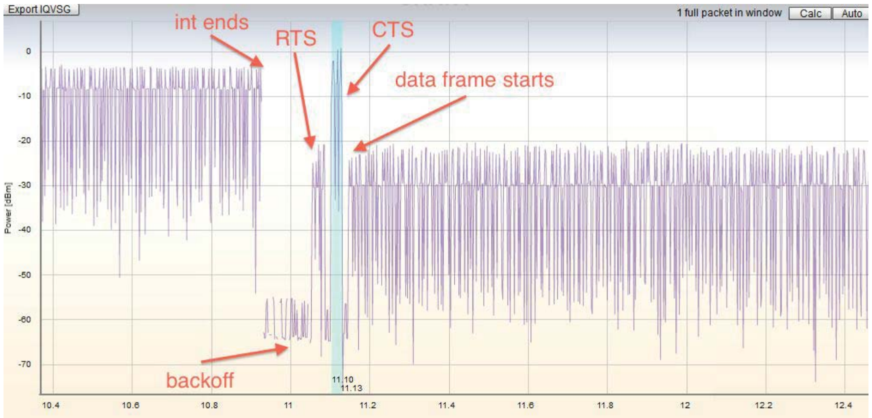

Note that the Wi-Fi devices have no knowledge of the LTE-U interference or its pattern, and thus once the medium is clear, Wi-Fi transmits frames in the normal manner, without any anticipation that LTE-U will begin transmission at some point. This point is critical to the following figure, which zooms in on the start of the LTE-U ON cycle.

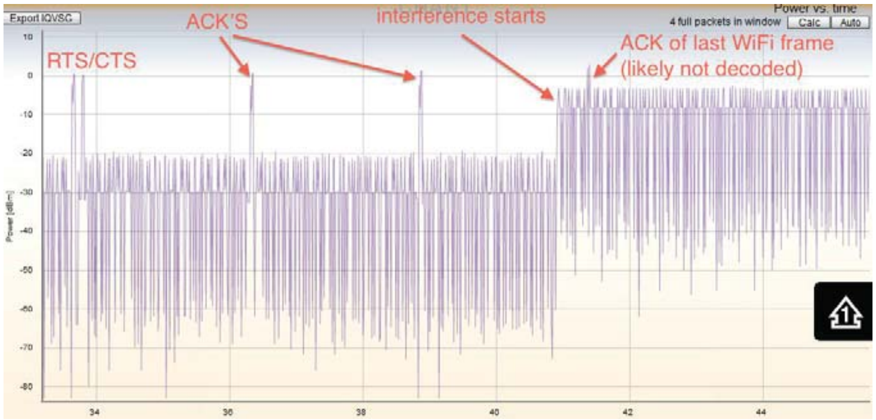

An RTS/CTS followed by 3 long frames is seen, the last of which begins at approximately 39 ms. LTE-U starts transmission at 41 ms, and thus interferes with the Wi-Fi frame that started at 39 ms. An ACK is sent in response to this Wi-Fi frame, but no other Wi-Fi transmission is made because the ED threshold is exceeded at each Wi-Fi device.[10] The high power interference from LTE-U corrupts this last Wi-Fi frame and likely also corrupts the subsequent ACK, and this turns out to be the critical reason why LTE-U duty cycling can be so harmful to Wi-Fi. From the perspective of the Wi-Fi transmitter (which is the client device, in this case), a frame error is observed[11] and it is generally not known if the error was due to interference or due to a change in propagation conditions.[12] Such a frame error occurs at the start of every LTE-U ON cycle,

---

[10] The Wi-Fi protocol dictates that an ACK be sent even if the medium is busy.

[11] An error is avoided only if both the data frame and the subsequent ACK are both received correctly. This is extremely unlikely in the presence of high power interference from LTE-U.

[12] Since the Wi-Fi client device was transmitting (and thus not receiving) at the instant when LTE-U began transmitting, the client device cannot differentiate whether the frame error was due to interference or some other reason, such as a change in the channel propagation conditions.

and this can in turn lead to *rate control* repeatedly reducing the transmitted rate and thus a throughput that rapidly spirals downward.

Rate control refers to the algorithm used by a Wi-Fi device to determine the transmitted rate of each transmitted frame, with the basic objective being to find the highest transmit rate that is supportable by the current link condition. Almost all Wi-Fi devices perform rate control based on the success/failure history of previous frames sent to the recipient, and a common behavior is that a sufficient number of frame errors triggers reduction of the rate.[13] Since LTE duty cycling causes a frame error at the start of *every* LTE-U ON cycle, a rate reduction can occur after a few such cycles are experienced, and since the errors continue, a further reduction may be seen afterwards, potentially leading to the transmission rate setting on a very low value. Further exacerbating this problem is the fact that lowering the rate increases frame duration (frame duration is the number of bytes divided by the rate). Longer frames means a smaller number of frames fit per LTE-U OFF cycle, and thus a higher percentage of frames are in error due to the start of LTE-U transmission, which can in turn lead to still further rate reduction.

The below figure displays the transmitted rate versus time for the client device; the left side shows an entire 15 second capture, while the right is a zoomed-in portion. During the first approximately 5000 ms, LTE-U is turned completely off to establish an interference-free baseline, and Wi-Fi reliably transmits at 260 Mbps. LTE-U begins transmitting with a 50% duty-cycle (30 ms on/30 ms off) at approximately 5000 ms, and then completely shuts off again at 13000 ms. The transmitted rate rapidly drops once the duty-cycled LTE-U is turned on and settles at a rate around 120 Mbps, but then rapidly jumps back to 260 Mbps when LTE-U is shut off at 13000 ms. The zoomed-in plot on the right shows the dynamics of this rate drop in more detail. Before LTE-U begins at approximately 4950 ms, a solid line of red markers is seen at 260 Mbps, indicating that Wi-Fi is continuously transmitting frames. Once LTE-U begins (at 4950 ms), short bursts of transmissions (with duration equal to the LTE-U OFF time) are seen at 260 Mbps; these bursts are because the Wi-Fi transmitter is turning off whenever the duty-cycled LTE-U is on. For the first 6 LTE-U periods, Wi-Fi largely holds a data rate of 260 Mbps.[14] But as discussed earlier, an error occurs at the end of each of those periods and this eventually triggers a reduction in rate, and this process continues until some intermediate settling is seen at a value slightly above 150 Mbps (the left figure shows the rate eventually settles around 120 Mbps). Instead of a throughput of 130 Mbps (half of the 260 Mbps rate observed for Wi-Fi in the absence of LTE-U), LTE-U operating at a 50% duty cycle reduces the Wi-Fi throughput to 60 Mbps (half of the 120 Mbps achieved by Wi-Fi when in the presence of duty-cycled LTE-U)—a factor of 2 smaller than expected.

---

[13] Rate control is a very vendor-specific aspect of Wi-Fi, and a wide variety of algorithms are seen. Some algorithms automatically reduce the rate of the subsequent frame whenever a frame error occurs, while others reduce rate only after multiple errors are observed within some time window.

[14] A few attempts with rates just under 200 Mbps are also seen, and these may be due to an automatic reduction of the rate of just the first frame (but not subsequent frames) after an error occurs.

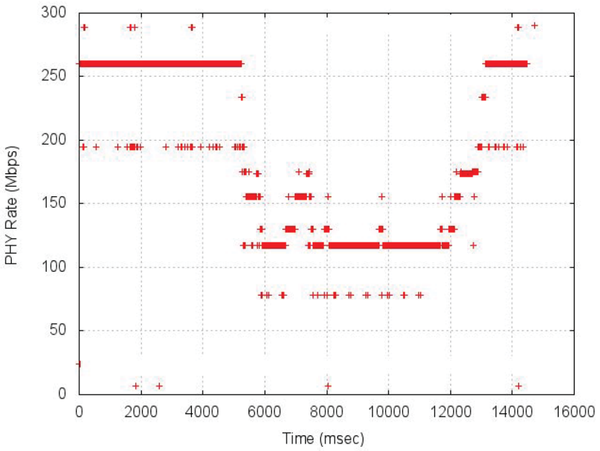

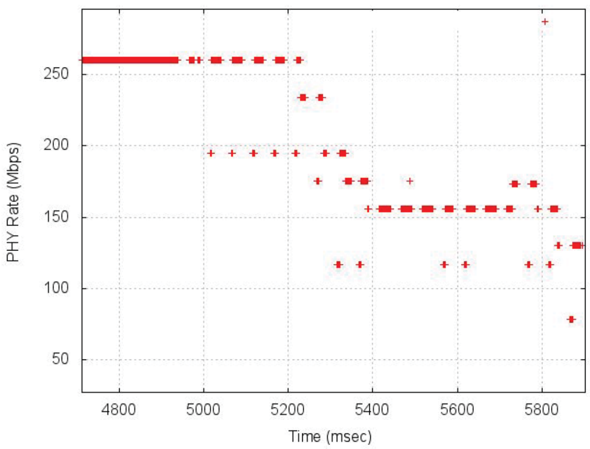

The percentage of Wi-Fi frames that are received erroneously due to the start of LTE-U transmission is a critical factor, because a higher percentage of frames received erroneously corresponds to a greater likelihood that rate control decreases the rate. As a result, short LTE-U OFF times lead to a high percentage of errors and thus can lead to poor throughput, whereas long OFF times can result in a low percentage of errors and thus better throughput.

This finding is confirmed in the below figure, which displays normalized Wi-Fi throughput versus LTE-U OFF time for a fixed duty cycle of 50%. Normalized throughput is defined as the Wi-Fi throughput in the presence of LTE-U divided by the Wi-Fi throughput in the absence of LTE-U. This quantity is computed by first running a throughput test without LTE-U interference, and then re-running the test with LTE-U turned on. Each curve corresponds to a test conducted for a specific pair of AP and client devices. (APs and client devices from different chipset vendors were included in testing.) For a 50% duty-cycle, effective timesharing corresponds to a normalized throughput of 0.5. Devices 1 and 2 effectively timeshare for the entire range of LTE-U OFF times (100 to 10 ms). On the other hand, timesharing is not effective for Devices 3 and 4, and an OFF time of 100 ms is required for these devices to achieve a normalized throughput close to 0.5 Finally, Device 5 is an example where timesharing is effective for OFF times of 50 ms or longer, but normalized throughput drops precipitously once the OFF time is reduced below 50 ms.

This diversity of results is representative of our general observations, which is that in some cases (i.e., for some propagation conditions and/or certain vendor devices) timesharing works effectively (i.e., rate control does not reduce the rate due to duty-cycled LTE-U) even for short/moderate LTE-U OFF times, while in other cases timesharing works effectively only for very long LTE-U OFF times (i.e., 100 ms, which is double LTE-U's maximum OFF time for a 50% duty-cycle) or not at all. Given the proprietary nature of rate control, variation in terms of how different devices react to duty-cycled LTE-U is expected, and is clearly observed. In addition, even small changes in the physical placement of Wi-Fi devices can significantly change timesharing efficiency. Sensitivity to physical placement and relative signal levels was observed for every device tested, meaning that no device that was tested always timeshared efficiently. For example, although Device 1 is shown to timeshare efficiently in the figure, that same device was observed to not timeshare efficiently in other configurations.

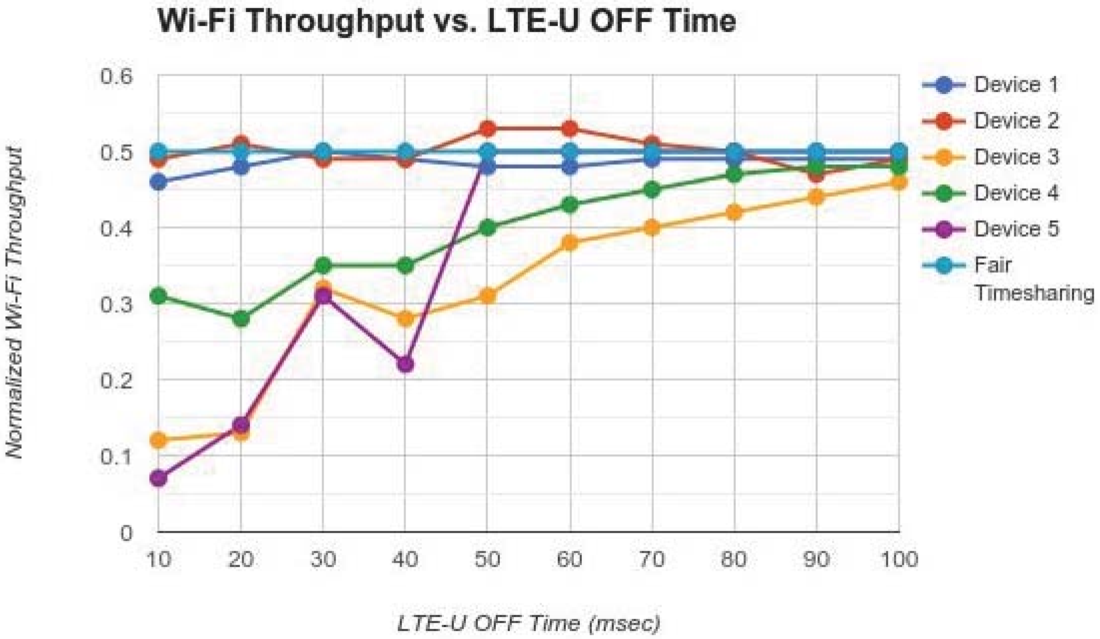

In summary, our testing shows that longer LTE-U OFF times reduce but do not necessarily eliminate the degradation in Wi-Fi throughput due to frame errors that arise from LTE-U duty-cycling. An OFF period of less than 20 ms is seen, across a variety of device types, to quite consistently lead to poor Wi-Fi throughput. OFF times of 50-100 ms lead to considerably better Wi-Fi throughput, although even with such OFF times, Wi-Fi degradation is often observed.[15] This behavior is an inherent consequence of LTE-U's failure to monitor the channel before beginning transmission. In addition, Wi-Fi hardware was not designed in anticipation of such interference, and hence does not necessarily respond well to it.

While long LTE-U OFF times are preferred from the perspective of Wi-Fi throughput, good delay and latency performance by Wi-Fi devices requires short LTE-U ON times. This is because Wi-Fi devices refrain from transmitting whenever LTE-U is transmitting, thereby introducing a specific delay pattern to all Wi-Fi traffic. This delay is quite consequential for high QoS traffic such as VoIP, video playback, video conferencing, and online/cloud-based video gaming, all of which tend to send short bursts of data every tens of ms. For example, a typical implementation of the ITU-T G.711 video encoder operates at 64 kbps and transmits a 160 byte packet every 20 ms. This pattern would clearly be interrupted by periodic LTE-U ON bursts, and would ultimately increase the end-to-end delay due to the additive nature of network latency. In a similar vein, video frames are often sent in bursts every 33 ms and would thus be similarly affected.

---

[15] Although results are shown here only for a 50% duty cycle, extensive testing with other duty cycles confirmed the dependence on the absolute LTE-U OFF time, regardless of the duty cycle.

The Wi-Fi traffic delay introduced by LTE-U also has a significant impact on beacons and power-save. Power-save is a vital feature for Wi-Fi clients, and client devices operating in power-save mode generally wake up only for beacons—which are periodically broadcasted and contain, in addition to general information about the transmitting AP, information telling each client if it has queued data waiting for it—and then go back to sleep whenever data is not waiting for the client. Beacon transmissions are made according to a fixed period (typically 102 ms), and power-save client devices sleep and wake up precisely according to such a periodic schedule. In the presence of duty-cycled LTE-U, however, any beacon transmission that is scheduled to be sent during the LTE-U ON time gets delayed until LTE-U completes its ON cycle, because Wi-Fi is prohibited from transmitting during the ON cycle due to the ED threshold.

A power-save device expects to wake up, stay in receive mode for a short period of time until it receives the beacon (i.e., a few ms), and then go back to sleep, but the delayed beacon means that such a device may have to stay awake for tens of ms in order to receive the beacon. In fact, the average time by which a Wi-Fi beacon is delayed due to duty-cycled LTE-U is, assuming a 50% duty cycle, the LTE-U ON time divided by 4.[16] So with a 50 ms ON and OFF time, for example, a beacon is delayed by an average of 12.5 ms. The direct implication of this is that a power-save client device has to, on average, stay awake in listen-mode for 12.5 ms every time it wakes up, far exceeding the few ms that the device is designed to stay awake in the absence of active data transmissions. Such a long wait undermines the principal benefit of power-save, which allows a device to remain actively connected while dramatically reducing the percentage of time it has to be in active listen-mode.

LTE-U delayed beacons may also cause some client devices to disassociate from the AP. This is because some client-side power-save implementations have a cap on the maximum time a device stays awake listening for a scheduled beacon, e.g., a client may go back to sleep if no beacon is received within 10 ms of it waking up. This leads to a missed beacon, and a sufficient number of missed beacons can lead the client to disassociate with the AP, meaning connectivity is lost.[17] A Wi-Fi AP acting as multiple, virtual access points—a situation that is now quite common[18]—can face additional issues because each of these virtual APs has its own beacon and scheduled time, and being forced off the air for long periods of time can lead to one

---

[16] For any LTE-U period and assuming a random (and uniformly distributed) offset between the start of the Wi-Fi beacon cycle and the LTE-U cycle, the scheduled Wi-Fi beacon instants are uniformly distributed within the LTE-U on/off cycle. Thus, half of Wi-Fi beacons see no delay at all, while the other half are scheduled to start during an LTE-U on cycle. The delay of those beacons which coincide with an LTE-U on cycle is uniformly distributed between 0 and the LTE-U ON time, and thus the expected delay of those beacons is the LTE-U ON time divided by 2. By the law of conditional expectation, the overall expected delay is LTE-U on time divided by 4.

[17] It is worth noting that the amount by which beacons are delayed follows a highly regular pattern, instead of a random delay per beacon. For example, with beacons sent every 102 ms and a 100 ms LTE-U period, if one beacon is delayed by 50 ms, the next will be delayed by 48 ms, and the one that follows by 46 ms, etc. This can lead to long consecutive strings of missed beacons for clients that give up on beacon reception after some pre-defined time, thus making dis-association even more likely than the average beacon delay time indicates.

[18] For example, most retail APs sold today allow for private and guest networks, both on the same channel but each with a separate identifier (SSID).

or more of these beacons being dropped from the AP's queue. This can lead to additional client disassociation.

It has been proposed to puncture long LTE-U ON periods with short 1-2 millisecond gaps in order to allow Wi-Fi to transmit delay-critical packets[19], such as beacons and VoIP frames, but a number of fundamental issues exist with this method. First, there is no guarantee that the intended delay-critical frames will actually be transmitted during the short gaps. Wi-Fi is only aware that the air has gone clear, but is not informed of the fact that LTE-U will resume transmission shortly afterwards, meaning that the efficacy of these gaps depends solely on Wi-Fi devices transmitting delay-critical frames before other frames. Wi-Fi APs typically do prioritize overdue beacons by moving them to the head of the queue (before any other queued frames), but a collision will occur whenever multiple APs each attempt to transmit beacons (once the air is clear, beacons are transmitted after a deterministic—rather than random—backoff time). VoIP and other QoS frames are more likely to be transmitted than best effort data, but the Wi-Fi protocol does not provide absolute priority and thus there is still a reasonable chance that best effort data is sent instead of high-QoS data. Second, and most importantly, the introduction of these short punctures can in fact exacerbate the rate control issues highlighted earlier. This is because the start of LTE-U transmission at the end of a short gap will interrupt whatever Wi-Fi frame is on the air, thereby increasing the frame error rate (compared to LTE-U operating without puncturing) and thus the likelihood of rate control reducing data rate.

Section 3.2: LTE-U Lacks Coexistence Mechanisms for Moderate LTE-U/Wi-Fi Interference (below -62 dBm).

This section focuses on scenarios where LTE-U interference is heard at a Wi-Fi device below the -62 dBm energy detect (ED) threshold (and vice versa) but appreciably above the noise floor. Although this corresponds to a lower interference level than in the previous sub-section, the interference generated by LTE-U can actually be more harmful to Wi-Fi when it is below rather than above -62 dBm. To frame this discussion, note that the effective thermal noise floor seen by a 20 MHz device is -96 dBm (assuming a 5 dB noise figure), and the received powers required to correctly demodulate the lowest and highest Wi-Fi rates (6 Mbps and 96.3 Mbps, respectively, when operating over 20 MHz) are approximately -92 and -60 dBm, respectively. This indicates that the -62 dBm ED level in fact corresponds to a very strong signal, and hints at how damaging interference that falls below the -62 dBm threshold can be. In addition, recall that a Wi-Fi device defers to other Wi-Fi transmissions at a power level of -82 dBm or lower, via the preamble detection mechanism described in Appendix 3; thus, this issue is indeed specific to LTE-U/Wi-Fi coexistence.

To quantify the impact of LTE-U interference at different power levels, we activated a Wi-Fi AP-client pair (running full-buffer downlink traffic via iperf) in the presence of interference from an LTE-U eNB emulator. The Wi-Fi link quality was held constant while the power of LTE-U was varied from low to high; in this particular experiment, the Wi-Fi client receives the

---

[19] Qualcomm, *Achieving Good Coexistence*, https://www.qualcomm.com/invention/research/projects/lte-unlicensed/lte-u-wi-fi-coexistence.

Wi-Fi AP at a power level of -57 dBm and achieves an interference-free throughput of 102 Mbps. Due to the fact that LTE-U Forum coexistence tests (see Appendix 1) are only defined for scenarios where the LTE-U eNB receives the Wi-Fi AP(s) at a level above -62 dBm, expectations are that LTE-U duty-cycling will not be performed when Wi-Fi is received below -62 dBm. But since behavior for below -62 dBm scenarios is not explicitly specified by the LTE-U Forum, the experiment was conducted with LTE-U continuously transmitting (i.e., 100% duty-cycle) as well as with a 33% duty-cycle (33 ms/67 ms ON/OFF time).

The figure below plots Wi-Fi throughput versus the received power of the LTE-U signal at the Wi-Fi AP. The devices were positioned such that the LTE-U interference was heard at the same level at both the Wi-Fi AP and client, and the LTE-U eNB and Wi-Fi AP receive each other at the same power level Physically speaking, the experiment corresponds to placing two Wi-Fi devices a fixed distance apart, and then moving the LTE-U eNB along the line bisecting the Wi-Fi AP/client path.

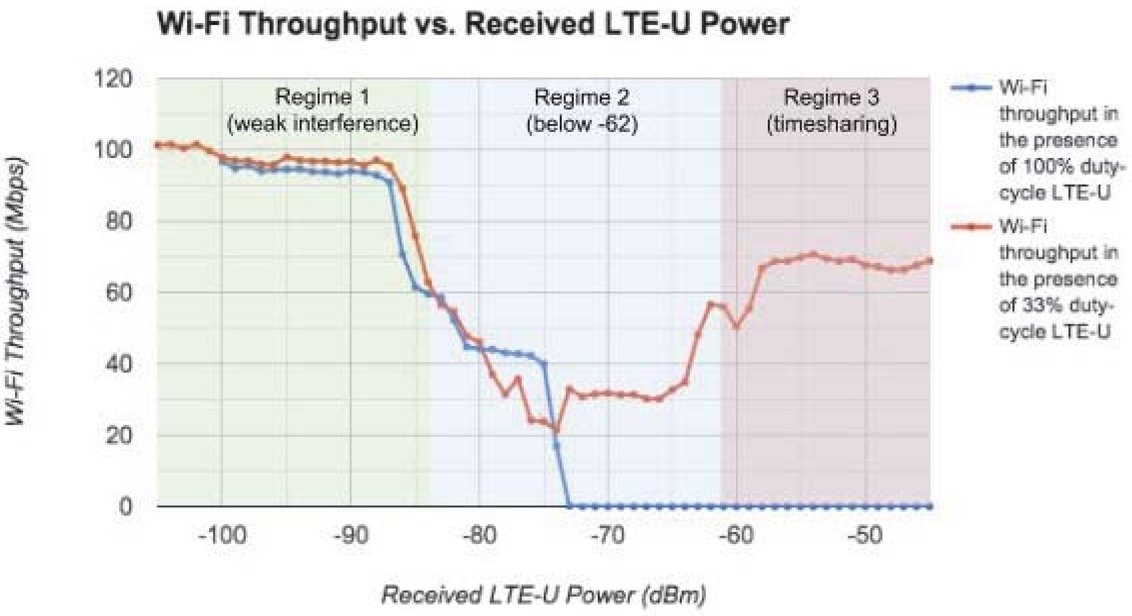

The blue curve, which shows Wi-Fi throughput in the presence of continuous LTE-U (100% duty-cycle), is best understood by tracing its behavior from left to right, i.e., low to high LTE-U power. At the far left, LTE-U is at or below the noise floor of the Wi-Fi device. Thus LTE-U has little or no effect on Wi-Fi performance, and the interference-free throughput (of 102 Mbps) is nearly or fully achieved. This section of the figure is labeled as Regime 1. As the LTE-U power is increased, Regime 2 is entered and a significant degradation in Wi-Fi throughput is seen. Continuous LTE-U transmission increases the effective noise floor of the Wi-Fi device, which translates to a reduced SINR (signal-to-interference-and-noise ratio) and thus a throughput that decreases monotonically with LTE-U power. Once continuous LTE-U is received at -72 dBm or higher, Wi-Fi is no longer able to maintain connectivity and throughput drops to zero.[20]

The red curve, which corresponds to Wi-Fi throughput in the presence of 33% duty-cycled LTE-U, is best understood by tracing it from right to left. In Regime 3 on the far right, LTE-U is received by Wi-Fi at a level above the -62 dBm ED threshold. In this particular case effective timesharing occurs[21], and Wi-Fi throughput is roughly two-thirds (the fraction of time LTE-U is off) the interference-free Wi-Fi throughput. Note that LTE-U power does not affect Wi-Fi throughput in Regime 3, because Wi-Fi transmits only when LTE-U is off.

Regime 2 occurs when the received power from LTE-U falls below Wi-Fi ED. This scenario results in the most significant degradation to Wi-Fi transmissions, and a very large throughput drop is observed as the interference from LTE-U falls below ED. This is because, first, the LTE-U interference is on the same order as the Wi-Fi signal, so that communication while LTE-U is on is either very limited or not possible at all, due to a low SINR (in this example, only if LTE-U is weaker than -72 dBm is communication possible). Second, the Wi-Fi devices are not aware of the duty-cycled nature of the interference. When the interference is above ED, the Wi-Fi devices are indirectly informed of the duty-cycle because Wi-Fi is forced, due to ED, to only transmit when LTE-U is off. However, when the interference is below ED, Wi-Fi devices will attempt to transmit when the LTE-U is on or off. In terms of rate control, one possibility is that the device settles on the highest rate that can be successfully transmitted in the presence of LTE-U, regardless of whether it is transmitting or not; this means that rate control will pick a low rate dictated by the SINR when LTE-U is on, regardless of the LTE-U duty cycle. This is because any rate that is appropriate for the LTE-U OFF cycle but is too high for the LTE-U ON cycle will essentially have a frame error rate equal to the LTE-U duty cycle, and such a high error rate is often adapted to by rate control algorithms. Thus, the rate control mechanism will slow down transmissions even when LTE-U is not on.[22] Another possibility is that the Wi-Fi transmitter is mostly able to limit its data transmissions to the LTE-U OFF cycle. This can happen, for example, if the Wi-Fi transmitter regularly uses RTS/CTS to precede data frames and those RTS/CTS exchanges fail whenever LTE-U is on, due to poor SINR in either direction of the Wi-Fi link. In such a case a flood of failed RTS/CTS exchanges occurs while LTE-U is on, and not until LTE-U turns off does the RTS/CTS exchange succeed, after which data frames are finally sent. Although potentially better than settling on a very low rate that works at all times, repeated RTS/CTS failures—which Wi-Fi devices typically interpret as an issue with the

---

[20] One might expect the Wi-Fi link to withstand a higher level of interference than -72 dBm, since the nominal SINR (recall that the Wi-Fi client device receives the Wi-Fi AP at a level of -57 dBm) with this interference power is 15 dB, which is higher than requirements for the lowest data rates. Note, however, that the client to AP SINR is lower than 15 dB due to the fact that Wi-Fi clients typically operate with a lower transmit power than APs, and also that different frequency-selectivity on the LTE-U to client versus on the LTE-U to AP link strongly affect the actual operational SINR.

[21] A large LTE-U OFF time was used for the 33% duty-cycle in an attempt to eliminate the rate control issues highlighted in Section 3.1, so that the experiment could focus on the unique aspects introduced from below-ED operation.

[22] Using a different transmitted rate when LTE-U is on versus when it is off is generally not feasible, because Wi-Fi rate control performs some averaging over the success/failure of recent frames and thus cannot react instantaneously to a change in channel conditions and/or interference.

link—very often lead to a reduction in the transmitted rate and thus this may not yield a better outcome than the first possibility. The key point is that both of these possibilities result in Wi-Fi throughput lower than that achieved by effective timesharing; the exact amount that throughput is lowered will depend on the vendor-specific rate adaptation algorithm implemented in the Wi-Fi devices. This is indeed the case in the figure, as the red curve levels out (in Regime 2) at a throughput which is approximately 60% lower than in Regime 3.

As the LTE-U interference power is further reduced, the 33% duty-cycle curve trends towards the 100% duty-cycle curve, and eventually Regime 1 is entered. Here, as well as in the transition between Regimes 1 and 2, it is observed that rate control selects a rate that works whether or not LTE-U is on, and consistent with this, the throughputs with 33% and 100% duty-cycled LTE-U are very similar.

A few higher-level observations can be made. First, if LTE-U employs duty-cycling only when above -62 dBm (in line with LTE-U Forum coexistence tests), then the Wi-Fi throughput would correspond to the 100% duty-cycle blue curve in Regimes 1 and 2, and the 33% duty-cycle curve in Regime 3. This would mean that the Wi-Fi link completely loses connectivity when LTE-U is heard at a level between -62 and -72 dBm, and not until LTE-U is received at -85 dBm or weaker does the Wi-Fi throughput recover to the effective timesharing level of 68 Mbps. Second, performing duty-cycling when LTE-U is received *below* Wi-Fi's ED level provides some gain relative to continuous LTE-U transmission in the right half of Regime 2, but duty-cycling when below ED is not nearly as effective as when LTE-U is received above ED. If below-ED duty-cycling was fully effective, then Wi-Fi throughput would not suffer any drop as the LTE-U power drops below ED because the power at which LTE-U is received should not affect throughput when timesharing works effectively (this is seen in Regime 3). But clearly this does not match the data shown here, nor in identical tests of other clients and APs.

In summary, LTE-U and Wi-Fi coexist relatively poorly when the devices hear each other below the -62 dBm ED threshold but above the thermal noise level. Duty-cycling improves the situation to some degree, but the technique is considerably less effective when employed below rather than above Wi-Fi's ED level. Although our discussion has focused on scenarios with symmetric received power, meaning that the LTE-U eNB and Wi-Fi AP hear each other at the same power level and the Wi-Fi client and AP receive the same power from the LTE-U eNB, it is important to note that asymmetric received power can make coexistence even more challenging. For example, if the Wi-Fi AP/client hear the LTE-U eNB above/below ED, respectively, this can create a hidden node scenario where only one of the two Wi-Fi devices can transmit data frames. This also brings up the question of whether or not LTE-U will perform duty cycling when the Wi-Fi client but not the AP is received above -62 dBm (such a situation is not covered by LTE-U Forum coexistence testing). Another possibility is that the LTE-U eNB receives a Wi-Fi AP below -62 dBm, while the Wi-Fi AP receives the LTE-U eNB above -62 dBm; if LTE-U does not duty-cycle, then the Wi-Fi AP is continuously held off the air because Wi-Fi's ED level is constantly exceeded.

Section 4: LAA Coexistence Challenges

While LTE-U is only targeted at regulatory regimes without LBT requirements, LAA is developing modifications to the LTE protocol so that the feature can also be used in regulatory regimes with an explicit LBT requirement, most prominently ETSI. (Relevant ETSI LBT requirements are detailed in Appendix 2.) Yet LBT is not by itself sufficient, for a variety of reasons, to address coexistence problems.

One of the most critical issues is how quickly a device begins transmitting after the medium is sensed as clear. Recall that the Wi-Fi protocol dictates that a device hold off transmission until the medium is clear (i.e., received power on the medium falls below a prescribed threshold), after which the device continues to hold off transmission for a short backoff time—and only if the medium is still clear after the backoff time does the Wi-Fi device transmit. If another unlicensed user transmits immediately after the medium goes clear (i.e., the device uses a zero backoff), then that device will always “win” access to the channel because its immediate transmission will shut out devices who perform a non-zero backoff. Unequal access also results when an LTE device uses a fixed non-zero backoff value or a different randomization of the backoff value such that the LTE device wins access to the medium considerably more often than a device using Wi-Fi’s random and increasing backoff distribution. Wi-Fi also employs an exponential backoff mechanism which doubles the backoff window whenever collisions occur.

As a result of these considerations, any technology employing LBT and a backoff technique different from Wi-Fi’s needs to very carefully consider coexistence between different types of devices. ETSI Options A and B each specify a backoff procedure that is different from Wi-Fi’s, and thus each can lead to inequitable channel access when compared to Wi-Fi devices. Option A uses an exponential backoff and so is similar in principle to Wi-Fi’s random backoff technique, but the exact backoff windows used are not the same as Wi-Fi’s and in many cases the windows are smaller than Wi-Fi’s. This can lead to Option A devices dominating channel access with respect to Wi-Fi users. Option B, on the other hand, uses a fixed backoff window (without any exponential increase), which can quite strongly tilt the share of access towards Option B devices and away from Wi-Fi in dense deployment scenarios where many users are contending for a channel.[23]

A second issue is that ETSI only requires LBT be performed to a threshold of -60 dBm, despite the fact that interference weaker than -60 dBm can be extremely harmful. This rule permits LAA devices to transmit so long as Wi-Fi or other signals are received at a level below -60 dBm. Deploying such devices would cause the coexistence problems discussed more fully in Section 3.2, where it was shown that a Wi-Fi link can suffer tremendous degradation—including complete loss of connectivity—from an always-on LTE-U transmitter heard slightly below the ED threshold. Generally speaking, an aggressive threshold such as -60 (or -62) dBm makes sense only in very dense deployment scenarios, where the received power

---

[23] Consider, for example, 3GPP document R1-151058, “Simulation results for Downlink Coexistence of LAA and Wi-Fi”, 3GPP TSG RAN WG1 #80 LAA ad hoc, for a detailed study of these issues that was completed for the 3GPP LAA process.

from all desired (non-interfering) sources is well above -60 dBm. Although dense deployment scenarios are very important, certainly not all deployments fit this description and designing coexistence rules that only work for dense deployments would unnecessarily limit the environments in which these unlicensed technologies are of use.

A further consideration has to do with transmission length. LTE, due to its basic protocol design and scheduled nature, generally transmits long frames (i.e., multiple ms), whereas a large percentage of Wi-Fi frames are sub-millisecond in duration. For this reason, equitable access to the medium, evaluated in terms of how often a technology is able to start a transmission, does not necessarily translate into equitable airtime.

The use of LBT could create the pathway for LAA to coexist gracefully with Wi-Fi and other unlicensed technologies, but only with very careful study of coexistence issues (coexistence studies are a part of LAA) in close collaboration with the broader unlicensed community, e.g., IEEE 802.11 and the Wi-Fi Alliance.

Section 5: Policy Questions

Sections 1 through 4 of this paper focus exclusively on technical evaluation of coexistence between LTE in unlicensed and Wi-Fi. In this section, we take a different tack and explore coexistence from a policy perspective. A diverse variety of technologies have managed to coexist successfully in the unlicensed bands. A distinctive aspect of LTE in unlicensed—as compared to other unlicensed technologies developed to date—is that it is a license-anchored system that operates simultaneously across licensed and unlicensed bands. Furthermore, LTE in unlicensed allows traffic to be moved dynamically, on a per-user and even on a per-traffic flow basis, across the licensed and unlicensed bands. This makes LTE in unlicensed substantially less sensitive to interference and collisions in the unlicensed band, because it is able to move traffic so quickly from the unlicensed band to the licensed band, in a very granular fashion, whenever congestion occurs in the unlicensed band. Purely unlicensed operations, by contrast, can fail entirely if there is interference in that spectrum. Reduced sensitivity to the conditions in the unlicensed bands significantly reduces the incentives that designers of LTE in unlicensed have to develop well-functioning coexistence mechanisms.

In the scenario where multiple unlicensed-only technologies operate in a band, the coexistence scenario is symmetric, because no technology has alternative, exclusive-use spectrum resources. If one technology acts in what is perceived as an overly-aggressive manner, for example transmitting without any regard to other nearby systems, other technologies might react (at a short and/or long time scale) aggressively in return. The threat of such mutual self-destruction deters overly aggressive behavior, and incentivizes each competing technology to compromise and develop mechanisms for reasonable coexistence. The evolution of Bluetooth and Wi-Fi to enable these two technologies to operate simultaneously in the 2.4 GHz band is an example of such.[24] In contrast, a license-anchored system that can move a set of users' traffic from unlicensed to licensed, while continuing to

---

[24] Bluetooth/Wi-Fi coexistence is achieved via adaptive frequency hopping by Bluetooth and the minimization of 40 MHz channelization by Wi-Fi via the "fat channel intolerant" bit and the requirement that Wi-Fi devices default to 20 MHz channels in the 2.4 GHz band.

serve other users on the unlicensed band, whenever that particular set of users experiences congestion in the unlicensed band, has much less to lose from a degradation in the condition of the unlicensed bands.

LTE-U's failure to define an effective coexistence mechanism for the scenario where LTE-U and Wi-Fi devices hear each other below the -62 dBm ED threshold provides one illustration of this concept. This coexistence scenario was discussed in detail in Section 3.2, where it was shown that a Wi-Fi link can suffer severe degradation when facing LTE-U interference, especially if LTE-U does not duty-cycle. An LTE-U link can in fact suffer the same consequence, because Wi-Fi devices only defer to energy detected above the -62 dBm. (The Wi-Fi system will attempt to be polite by also searching for lower power Wi-Fi preambles, but an LTE-U device does not transmit a Wi-Fi preamble.) At first glance, this appears to be the sort of mutually destructive scenario that equally harms both parties. However, the nature of LTE-U's license-anchored operation tips the scale of this encounter quite strongly in the favor of LTE-U: the LTE system can fall back to its licensed, exclusive-use resources if it does not prevail in a conflict with the Wi-Fi system. The Wi-Fi system, though, cannot operate at all in the presence of major interference.

This argument is further illustrated in the below figure, which depicts an LTE-U eNB and a Wi-Fi AP in the vicinity of each other, but with each device hearing the other at a power level below -62 dBm. Assuming the eNB and AP are both continuously transmitting, the devices in between the eNB and the AP experience a poor SINR (signal-to-interference-and-noise ratio), because these devices are nearly equidistant to their signalling and interfering sources and thus receive nearly equal signal and interference power. Such poor SINR devices are indicated by purple triangles (Wi-Fi clients) and red circles (LTE clients). On the other hand, the blue triangles and circles are devices which experience a moderate or good SINR, because each of these devices is considerably closer to its respective signalling source than to its interfering source. In the context of the plot of Wi-Fi throughput versus power received from LTE-U in Section 3.2, the poor SINR nodes essentially correspond to Regime 2, where interference is below ED but leads to very poor or even zero throughput, while the moderate/high SINR nodes are in Regime 3, where the interference is weak compared to signal.

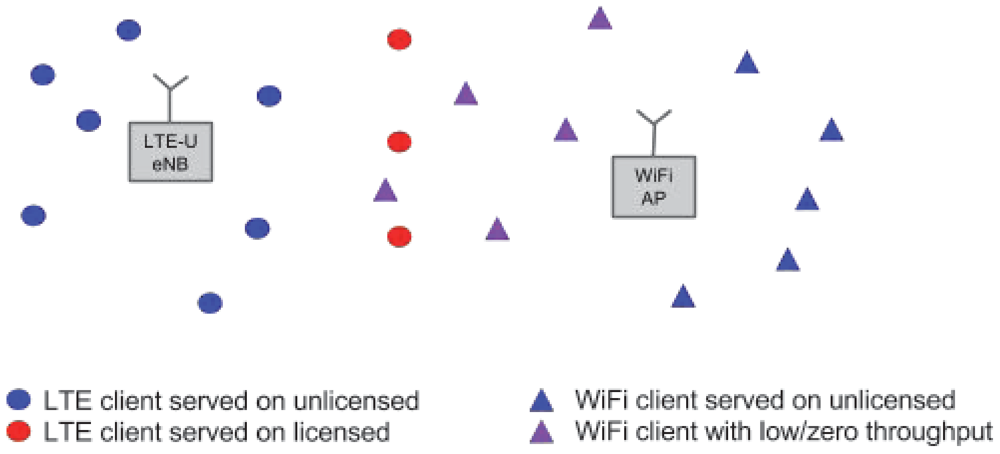

Due to the dynamic and flexible frequency-domain scheduling allowed by orthogonal frequency domain multiple-access (OFDMA) and carrier aggregation, it is straightforward for the LTE-U eNB to serve the LTE clients with a poor SINR in the unlicensed band (red circles) entirely on the licensed band (where there is no interference from Wi-Fi), while continuing to transmit data to the LTE clients with a good unlicensed-band SINR (blue circles) on the unlicensed band. However, Wi-Fi clients that are experiencing a poor SINR due to strong LTE-U interference have no such fall-back available. In fact, if the SINR is negative (in dB) or just slightly positive, it may not be possible for the client to decode even the lowest Wi-Fi rate—meaning that connectivity is entirely lost. The loss-of-connectivity scenario is worst-case but is not at all unrealistic, as made clear in Section 3.2. And even in a less pessimistic scenario where the Wi-Fi link remains marginally functional, LTE-U can still lead to a dramatic reduction in Wi-Fi transmission rates for clients receiving strong LTE-U interference.

The end result is that LTE-U suffers some degradation due to co-channel Wi-Fi interference but LTE overall can still operate quite successfully. Yet the Wi-Fi system, which does not have an interference-free licensed band to fall back on, is essentially crippled, because a substantial percentage of its devices either lose connectivity entirely or have to operate at a dramatically reduced data rate.

While the above picture shows LTE clients positioned closer to the eNB than to the Wi-Fi AP and vice versa for Wi-Fi clients, in reality we expect sticky association, where LTE/Wi-Fi clients must associate to the LTE-U/Wi-Fi APs, respectively, regardless of which base station device they are closer to. This means that an even higher percentage of devices will be in a disadvantaged interference scenario—LTE clients might be close to the Wi-Fi AP and a relatively far distance from their serving eNB, which must in turn increase its transmit power towards those clients—making Wi-Fi even more susceptible to being crowded out.

In short, we identify as potentially significant the fact that LTE-U operators have a unilateral, licensed-spectrum, solution to avoid the ill-effects of conflicting uses of unlicensed bands, while Wi-Fi operators must rely on bilateral or multilateral solutions.

Section 6: Conclusions

Our experimental evaluation of LTE-U, the proprietary version of LTE in unlicensed being developed for countries without LBT requirements, found problems with LTE-U's duty-cycled approach to sharing with co-channel Wi-Fi devices, as well as with LTE-U's lack of an effective coexistence mechanism when Wi-Fi is received at a power level below -62 dBm. While the above findings were made in terms of Wi-Fi throughput degradation, the potentially large impact of LTE-U on Wi-Fi latency merits considerable additional study; for example, LTE-U's duty-cycle almost inevitably will cause issues with periodic beacons sent by Wi-Fi APs, which can in turn seriously impact power consumption at Wi-Fi client devices. The additional latency introduced by duty-cycled LTE-U needs to also be studied in the context of new very latency-intolerant services such as real-time gaming. Another issue that deserves attention is how LTE-U's channel selection process (i.e., determining which portion of the 5 GHz unlicensed band to operate in) will interact with automatic channel selection algorithms in deployed Wi-Fi equipment. Given the huge diversity of deployed Wi-Fi equipment, it is critical that coexistence testing be performed against a suitably comprehensive set of AP and client. Scenarios with asymmetric received powers at different devices (LTE-U eNB and Wi-Fi) were briefly mentioned in Section 3.2, but also deserve considerable additional attention. Although our focus has been on Wi-Fi, other unlicensed-only technologies also operate in the 5 GHz unlicensed bands and LTE-U needs to be able to coexist with those as well. The LTE-U Forum does not define any technology-agnostic or non-Wi-Fi coexistence tests, so this issue requires further consideration.

LAA, which is the 3GPP standardization body's incorporation of LTE in unlicensed into LTE-Release 13, is adopting LBT as required by ETSI standards. Fair coexistence with Wi-Fi may be possible if LAA implements LBT effectively, meaning that access to the channel and the power level at which transmissions are deferred are both made equitable relative to that of Wi-Fi devices such that fair sharing with Wi-Fi is achieved. But only with close collaboration with all stakeholders, especially IEEE 802.11 and the Wi-Fi Alliance, can this outcome occur.

LTE in unlicensed represents yet another innovative technology unleashed by accessibility to unlicensed spectrum. But given the vast scale and tremendous importance of technologies currently operating in the unlicensed bands, Wi-Fi in particular, considerable further study is required to fully understand how LTE in unlicensed will coexist with those technologies.

Appendices

Appendix 1: LTE-U

The LTE-U carrier aggregation profiles specify a single component carrier in a licensed band acting as the PCell, and one or two 20 MHz component carriers in the 5 GHz unlicensed band; thus, up to 20 MHz in licensed spectrum and up to 40 MHz in the 5 GHz unlicensed band can be aggregated. The PCell supports downlink and uplink traffic, but the component carriers are downlink-only (Supplemental Downlink, in LTE terminology). Component carriers (each 20 MHz) in the unlicensed 5 GHz band are defined for the UNII-1 (5150-5250 MHz) and UNII-3 (5725-5825 MHz) bands, and have the same center frequencies as Wi-Fi channels 36, 40, 44, 48, 149,153,157,161, and 165, plus an additional channel at 5160 MHz. Channels are not defined within the UNII-2 (5250-5725 MHz) band.[25]

In terms of coexistence with existing Wi-Fi devices in the 5 GHz band and other LTE-U eNBs, the Forum provides a high-level definition of the following three coexistence mechanisms:

1. Channel Selection: an LTE-U eNB should attempt to choose a clean channel in which no other LTE-U eNB or Wi-Fi AP is operating.
2. Duty Cycling [Secondary Cell DTX]: If a clean channel cannot be found, then the SCell is duty cycled to share with co-channel Wi-Fi and LTE-U. It is recommended that the duty-cycling parameters be adaptively tuned, e.g., in response to Wi-Fi usage.
3. Opportunistic SCell Switch-OFF: The SCell is completely turned off when the LTE traffic demand can be met on the licensed band PCell.

Details of the above mechanisms are left entirely to manufacturers, but the Forum's coexistence specification document[26] does define the following explicit coexistence tests:

1. Channel Selection:
    a. If two channels are available and the first is unoccupied while the second is occupied by either a Wi-Fi AP or another LTE-U eNB that is heard above -62 dBm at the LTE-U eNB under test, check that the LTE-U eNB under test picks the first channel.
    b. If two channels are available, the first by an LTE-U eNB of the same operator and the second with an LTE-U eNB of a different operator, check that a new LTE-U eNB selects the first channel.
2. Duty-Cycling with 1 or 2 Co-Channel Interferers: If an LTE-U eNB is operating on a channel occupied by one single full-buffer Wi-Fi AP, two full-buffer Wi-Fi APs,

---

[25] LTE-U Forum, *UE Minimum Requirements for LTE-U SDL* V1.0 § 5 (2015-02), http://www.lteuforum.org/uploads/3/5/6/8/3568127/lte-u_forum_ue_minimum_requirements_for_lte-u_sdl_v1.0.pdf.

[26] LTE-U Forum, *LTE-U SDL Coexistence Specifications* V 1.0 (2015-02), http://www.lteuforum.org/uploads/3/5/6/8/3568127/lte-u_forum_lte-u_sdl_coexistence_specifications_v1.0.pdf.

or one full-buffer Wi-Fi AP and one full-buffer inter-operator LTE-U eNB, then check that the LTE-U eNB under test uses a duty cycle no higher than 50%/33%/33%, respectively, and that the maximum ON time of the duty cycle is no longer than 50 ms.[27]
    a. Rel. 12 Discovery transmissions made during the LTE-U off period are not accounted for in the duty cycle computation.
3. Opportunistic SCell Switch-OFF: If an LTE-U eNB is associated with only a single LTE-U client and then that client is either removed from the coverage area or traffic to that client is stopped, then the SCell is checked to be in the OFF state within X seconds, where X is a manufacturer-declared time.

The tests are defined only for scenarios in which interfering devices are heard above -62 dBm. The duty-cycling tests are conducted with full-buffer Wi-Fi AP(s), and the LTE-U Forum does recommend the use of adaptive duty cycling based on observed channel usage. This implies that a duty-cycle more aggressive than specified in the tests (i.e., 50% when sharing with 1 device, and 33% when sharing with two) can be used by an LTE-U eNB when co-channel to Wi-Fi AP(s) that are observed to not be fully loaded.

Appendix 2: ETSI LBT Requirements

ETSI requires devices operating in the 5 GHz unlicensed band to use LBT. Most relevant here are ETSI's requirement for Load Based Equipment, which is defined as "equipment where the transmit/receive structure is not fixed in time but demand driven".[28] ETSI lays out three LBT options for Load Based Equipment, one of which is to use the exact LBT rules defined by the Wi-Fi protocol.[29] The other two options require an energy detect threshold of -73 dBm/MHz (-60 dBm for 20 MHz operation) and two different methods for performing random backoff, referred to as Options A and B. Option A allows a maximum transmission duration of 10 ms and has an exponentially increasing backoff window. Option B uses a fixed backoff window, which is fixed by the manufacturer of a device and is directly tied to the maximum duration of a transmission (i.e., a device that wishes to use a short backoff window is also limited to short transmission durations). Options A and B share some principles with Wi-Fi, but differ in important manners as detailed in Section 4.

Appendix 3: Wi-Fi Medium-Access Control Protocol

A complete description of the Wi-Fi medium-access control (MAC) protocol is far beyond the scope of this paper, so here we attempt to provide a high-level summary of the protocol,

---

[27] Full-buffer refers to a Wi-Fi AP that has a constant stream of traffic to transmit, as opposed to intermittent/bursty transmission by the AP. This is an important distinction to make, because the LTE-U Forum documentation recommends an adaptive duty cycle, possibly based on channel usage - indicating that, for example, a duty cycle larger than 50% can be used if co-channel to a single Wi-Fi AP without a full buffer.

[28] ETSI EN 301 893 v1.8.1

[29] This refers to the energy detect threshold (-62 dBm) and the clear channel assessment procedures defined in the IEEE 802.11 specification, which are summarized in Appendix 3.

with an emphasis on aspects related to the coexistence discussion underway. Interested readers can find a full description in the IEEE 802.11 standard. The Wi-Fi MAC protocol uses LBT, i.e., wait until clear, and exponential random backoff, and can be described in three parts:

1. Step 1 (Wait until clear): A Wi-Fi device that wishes to transmits senses the medium until it is clear. The medium is clear if neither of the following two conditions hold:
    1.1. Energy Detect (ED): Received power at the device exceeds -62 dBm.
    1.2. Preamble Detect: If a Wi-Fi preamble is decoded, the channel is busy for the duration of the corresponding Wi-Fi frame. The IEEE 802.11 specification requires that Wi-Fi preambles be decodable down to a power level of -82 dBm, although state-of-the-art Wi-Fi devices decode to a level of -90 dBm or lower.[30]
2. Step 2 (Random backoff): Once step 1 is complete, the device selects a short random backoff time. If the channel remains clear for the duration of the backoff time, then the device will begin transmission.
    2.1. The random backoff time is chosen according to X + N * 9 microsec , where N is a randomly chosen integer between 0 and the parameter CW, and X is the minimum backoff time. CW, which is short for contention window, is an adaptive (see Step 3) per-device parameter, with a default value of 15. The value of X does depend on the traffic type, but 27 microseconds is a typical value.
3. Step 3 (Exponential increase of backoff): If a collision occurs, each involved device doubles its current CW value. A device reduces its CW to its minimum value (typically 15) once a packet is successfully transmitted.

Step 1 dictates that once a device has started transmitting, all other devices within hearing range should remain silent until that initial transmission is completed. After the initial transmission has completed, Step 2 randomizes the time at which the devices within hearing range can attempt transmission. This makes it likely that a single device begins transmitting before any other device, leading other devices to mark the medium as not clear and thus preventing a collision. A collision still can occur if two devices pick exactly the same backoff time, so Step 3 enlarges the window whenever a collision occurs (from the perspective of the transmitter device, a collision maps to transmission of a frame without reception of an acknowledgement) to reduce the probability of a subsequent such collision.

The size of the “hearing range” around an active transmission that is kept quiet via Step 1 is critical, since this determines the range at which devices can interfere with one another. In this regard, it is important to note that a Wi-Fi device defers to another Wi-Fi device (via preamble detect in Step 1.2) at a power level at least 20 dB lower than the -62 dBm ED threshold (Step 1.1). This translates into a dramatic difference in distance: the -62 dBm

---

[30] A Wi-Fi frame consists of a preamble followed by a data payload. Encoded within the preamble (which ranges from 20 to 52 microseconds, depending upon precisely which 802.11 version is being used) is basic control information describing the frame, such as the length (which can be up to 5.484 ms) and communication parameters (e.g., the data rate) of the payload section. The preamble is transmitted using the most robust (i.e., lowest) data rate in order to broadcast the presence (and length) of the frame to the largest possible neighborhood—precisely to keep nearby devices quiet per Step I.

threshold corresponds to 50 meters of distance, whereas -82 dBm corresponds to 200 meters.[31] This is particularly important in light of the fact that ETSI LBT specifications require energy detection to a threshold of -60 dBm (quite similar to Wi-Fi’s threshold) but do not have an additional, more stringent requirement such as Wi-Fi’s preamble detection threshold of -82 dBm.

---

[31] These numbers assume 30 dBm transmit power, as allowed in UNII-1 and UNII-3 in the United States, and the 802.11n channel D NLOS propagation model (free space for the first 10 meters; path loss exponent of 3.5 beyond 10 meters).