The Promise of WiFi in the 6 GHz Band

A Technical Paper prepared for SCTE•ISBE by

J.R. Flesch
Director, Advanced Technology
Commscope
3871 Lakefield Drive, Suwanee, GA 30024
678 473 8340
Jr.flesch@commscope.com

Charles Cheevers
CTO/CPE
Commscope
3871 Lakefield Drive, Suwanee, GA 30024
678 473 8507
Charles.cheevers@commscope.com

Kurt Lumbatis, Commscope

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

The 5.925-7.125 GHz band (colloquially “6 GHz band”) represents an immense opportunity for indoor WiFi to fully adopt the promise of WiFi6 in a green space environment and clear out the channel access baggage and heterogeneous technical epoch mix accumulated during the more or less organic growth of unlicensed, contention-based wireless services in the 2.4 and 5 GHz bands. In exploiting this clean break, it avoids disrupting the existing population of devices and their present state of interoperability (however suboptimal that may be). Spectrum leverages associated with multiple-user OFDMA, multiple-user MIMO and BSS coloring have the ability to promote low-latency spectrum scheduling, improved link margins and topographical channel re-use which will go a long way towards resolving the potentially thorny CCI environment represented by multiple-AP, dense client device deployments. Additionally, the wealth of new spectral piping available at 6 GHz, of itself, may provide all the solution required to wirelessly backbone data hauling in the home between WAN attachment and an opportunistically-placed AP/hub/extender, resulting in reliable (virtually OOB) trunking hauls between access points which enable whole-home LAN bitrates sufficient to meet anticipated WAN bulk connectivity budgets as these inflate via either DOCSIS or 5G mechanisms. Blanket WiFi coverage of multi-Gbps (as a services ensemble) ought to be achievable given the power, BW, link budget and spectral efficiencies available within the service radii posed by indoor residential environments.

The Exciting Promise of Indoor 6 GHz

1. Motivation for Exploit of 6 GHz as Unlicensed Spectrum Relief

1.1. Overcrowding in legacy unlicensed spectrum

The following figure details the spectral occupation associated with existing ISM and U-NII bands at 2.4 and 5 GHz which host 802.11-based wireless traffic across at least three technical specification epochs of that standard. MAC differentials across these epochs contribute to the access pathology by collapsing throughput to least-common-denominator type of wireless medium exploit in cases where heterogeneous client populations comprising some number of older legacy devices compete for airtime from the access point. Even in the cases where relatively high bitrate streaming traffic is shunted off of the 2.4 GHz to the 5 GHz band, crowding in both pieces of spectrum is becoming everyday more commonplace.

Figure 1 – Spectral Crowding @ 2.4 and 5 GHz for MDU cases

And the problem is exacerbated by IoT radios exploiting the 2.4 GHz ISM band to connect constrained end devices (potentially yielding status-critical telemetry in small packet traffic) to IoT mesh hubs (as discrete CPE or WiFi extender adjunct stackware) for ultimate backhaul over the 802.11 network. Home security and aging in place services represent two such IoT applications which can ill afford excessive attachment latencies (or worse – lost data). The upshot is that offset carriage spectrum needs to be mined in order to free up IoT access (for those services' several NFC MACs operating at 2.4 GHz) by moving 802.11 traffic away from that highly contentious band.

1.2. MDU and near-neighbor ingress

Allowable legacy wireless power levels conspire to recruit ingressing, unwanted interferers in the case of near neighbors – or even more problematically, MDU structures. While clever amelioration techniques like EasyMesh can identify problem channel competition and provision better-case utilization of available spectrum, such techniques are rendered less effective by overlays of wireless networks representing disparate control authorities (whose closed circuit loop dynamics can conspire to orchestrate chaotic thrash in the mixed environment). Migration to more common adoption of these type of higher stack layer controls will help – but not as much as reserving pristine new spectrum and reserving it for exploit by devices compliant to only the most recent MAC initiatives implemented in WiFi6.

1.3. The insatiable demand

To compound the spectral crowding, in-home wireless bitrate appetite is only increasing. Note the implacable demand for ever more device connectivity expected in the immediate future, as witnessed by the accompanying figure:

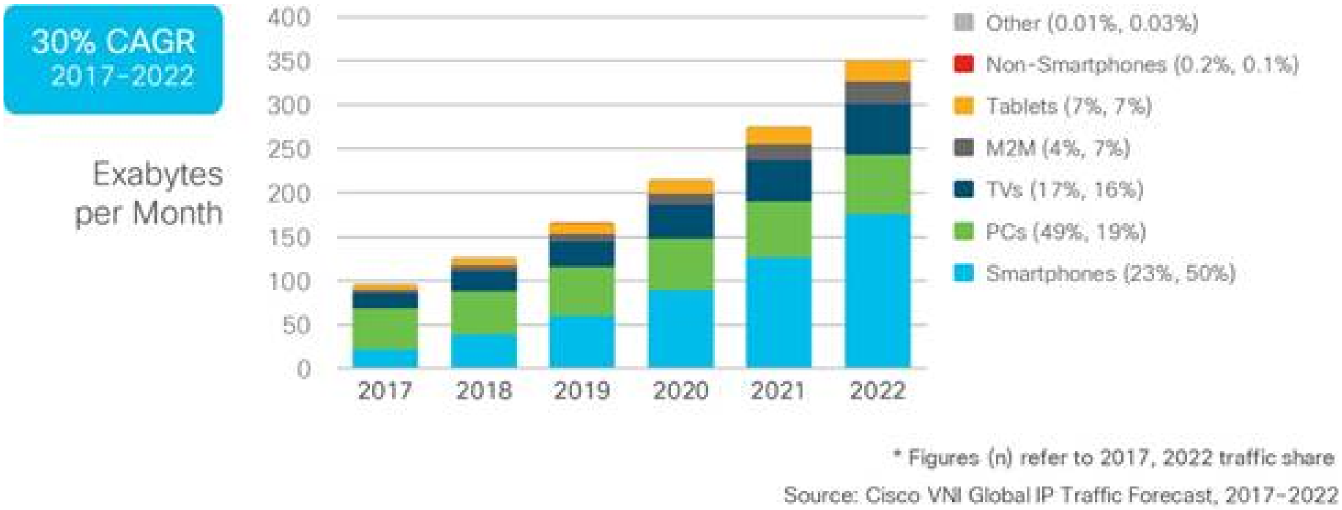

Source: Cisco VNI Global IP Traffic Forecast, 2017-2022

Figure 2 – IP traffic demand expectations

And this is for native 802.11 traffic. As alluded to above, IoT home device use is on the rise as well and this will inject further traffic into the spectrum, particularly at 2.4 GHz (and with Zigbee, Bluetooth and Thread MAC behavior which may not be at all WiFi-aware):

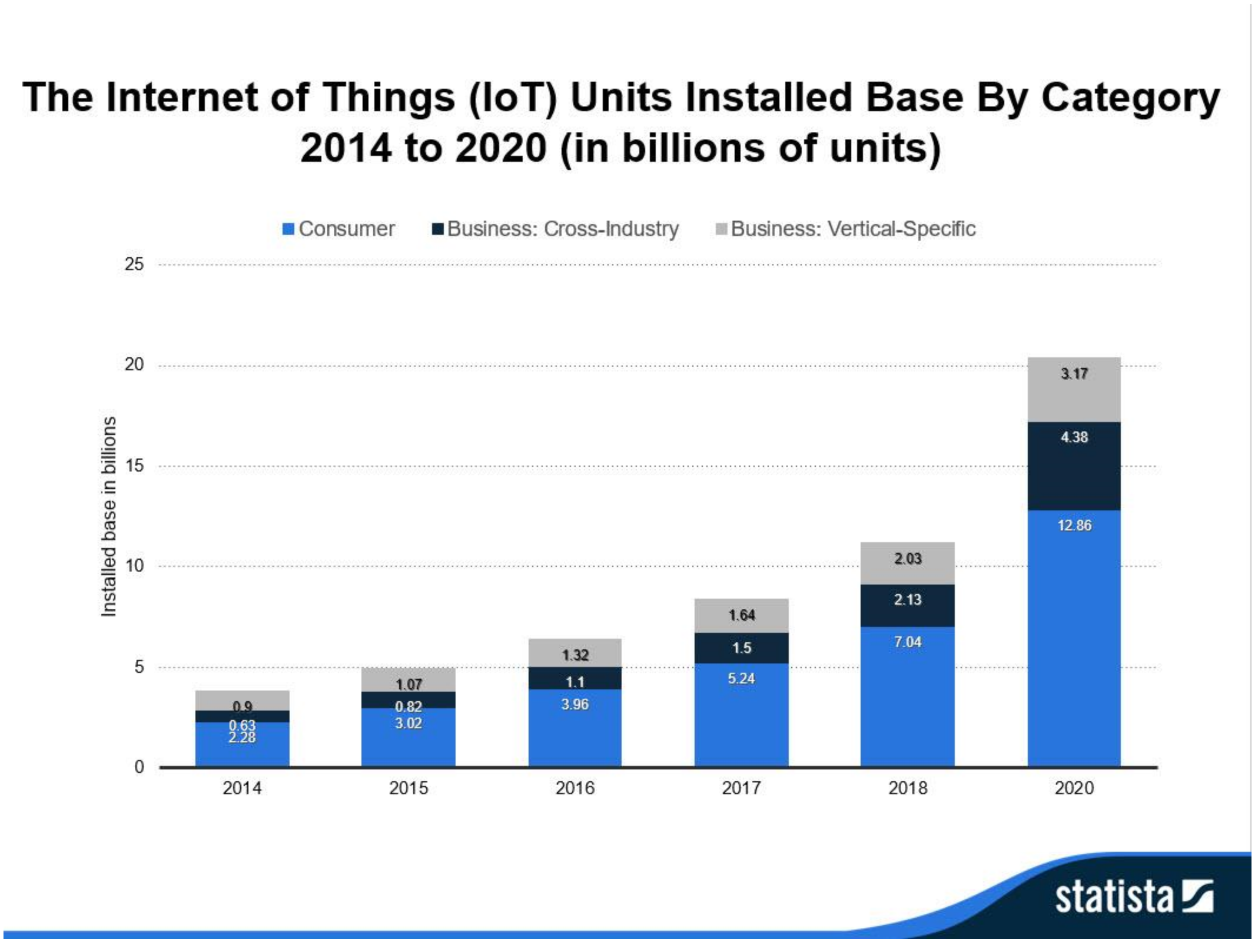

Figure 3 – Consumer adoption of wireless IoT devices by year (blue component)

1.4. Coverage gaps of the WAN gateway

Multilevel single family dwellings represent a single AP wireless coverage challenge once floorplans exceed 2000 square feet or so; this is a function of service radius (set by losses to both in-air link endpoint distances and accumulated solid surface transitions of walls and floors), increasing number of client devices competing for shared air time and the common restriction that wireline WAN services are usually introduced to the home at an external wall. This aperture typically wastes half of the isotropic radiated antenna pattern of the gateway located just at the interior feed point of the WAN and requires a service radius extender somewhere toward the middle of the home's interior to support wireless backhaul (or the co-option of existing – or the pulling of new – wireline to provide front- and backhaul support to the extender). Absent the extender, whole sections of the home furthest from the gateway would be effectively blanked from network attachment (certainly from bitrate service which could support multiple large-format video streaming applications, say).

Wireline trunk upgrades to the home, however, are not the stuff of self-installation and have the capacity to generate significant consumer ill will – the common feeling being that the ISP has failed in its duty to provide wireless coverage as was expected. And home infrastructure upgrade/maintenance costs are seldom properly anticipated and never tolerated well. MoCA, Ethernet and powerline co-option techniques do offer the promise of less-interference-susceptible trunking to the extender – but not necessarily to carrier-grade reliability levels (or without – excepting Ethernet – the use of two mirror end-of-link devices to transceive and transcode the wireline modulation scheme employed).

Meanwhile, wireless trunking to the extender to/from the WAN GW, while facilitating the buy/self-install services remediation paradigm, suffers from "small print description" reliability issues (i.e., extender type may not represent the correct solution to the given coverage problem). This is because the extender can problematically cannibalize channel bandwidth and/or airtime from client device usage for applications in order to sustain the ultimate WAN connectivity across one or more hops of the whole in-home network. As such, extenders with insufficient bandwidth or older MAC technology might actually worsen the whole-home service experience. Triband extenders (meaning 2.4/5Lo/5Hi) and WiFi5 bring some interim relief for this proposition (by differential exploit of the 5Lo and 5Hi bands for front- and backhaul, say), but absent the WiFi6/6 GHz proposition cannot endow the wireless home with the ensemble bitrate support needed to meet coming wireline and 5G fixed wireless WAN speeds.

2. Application of the 6 GHz Remedy

2.1. Prospective performance

As mentioned above, if just a reliable high-capacity wireless trunk over new spectrum can be placed between GW and extender, both legacy and new WiFi devices will experience additional connection capacity (the former via the recovery of spectrum lost to TriBand or other trunking and the latter, via exploit of the newly available BW). In initial deployments, the WiFi6 MAC upgrades associated with MU-MIMO, MU-OFDMA and BSS coloring will not even need to be invoked to see immediate, massive performance gains; even at low power (250 mW EIRP), the constrained dimensions of an average indoor floorplan (~2600 sq ft) suggest accumulated link losses should not force the negotiated MCS to drop below midgrade levels at worst case. The upshot of this observation is that spectral efficiency should remain high over relatively massive WiFi bandwidths (up to160 MHz – or more if the FCC reserves additional U-NII bands for low power unlicensed service).

To evaluate these expectations for indoor WiFi coverage utilizing the 6 GHz spectrum, the Arris/Commscope WiFi test house (a trilevel "average" US home) was instrumented. Perspective views and floorplans follow:

Figure 4 – WiFi test house, front view

Figure 5 – WiFi test house, rear view

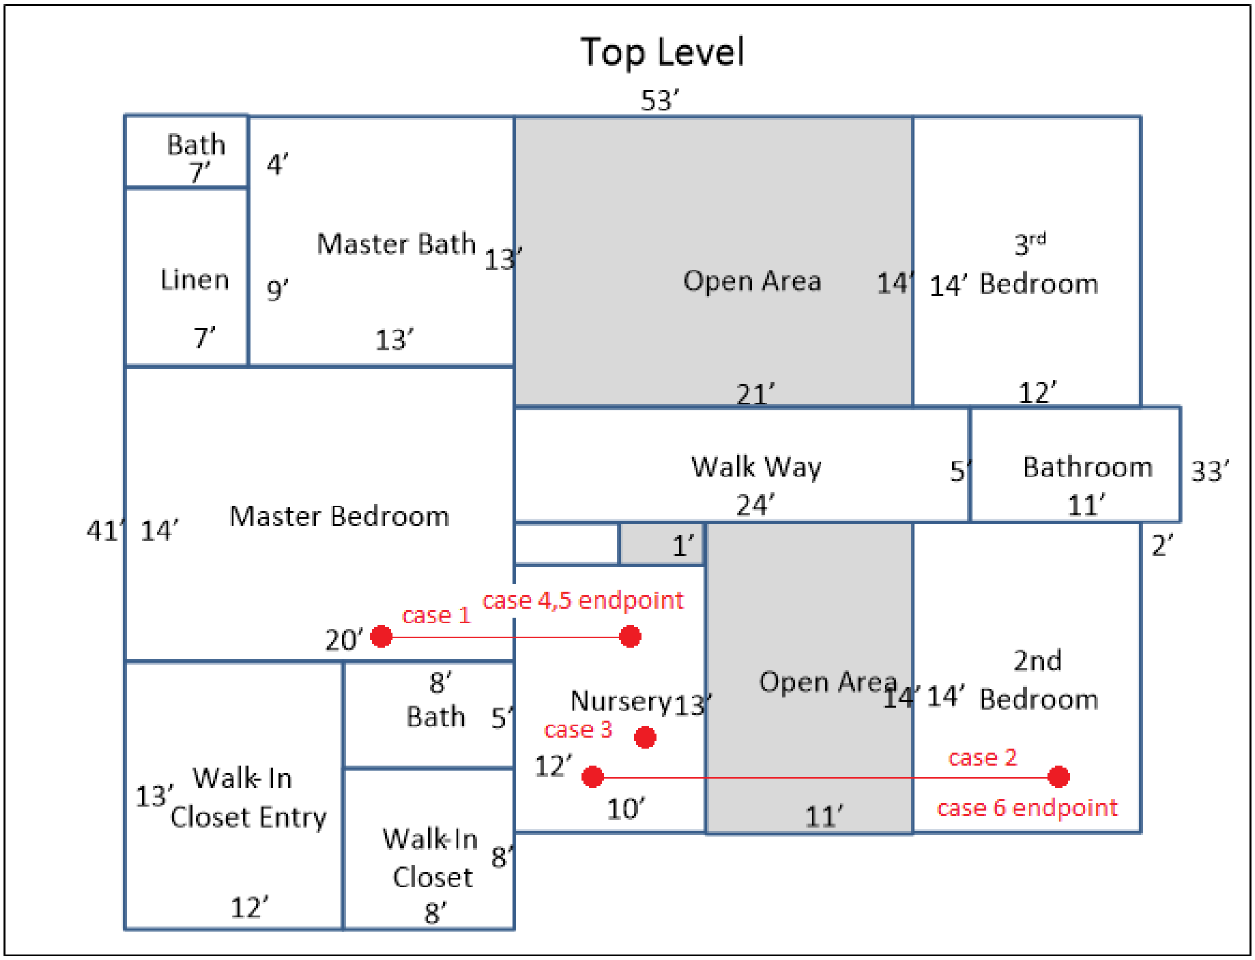

Figure 6 – WiFi test house, top level floorplan with test cases

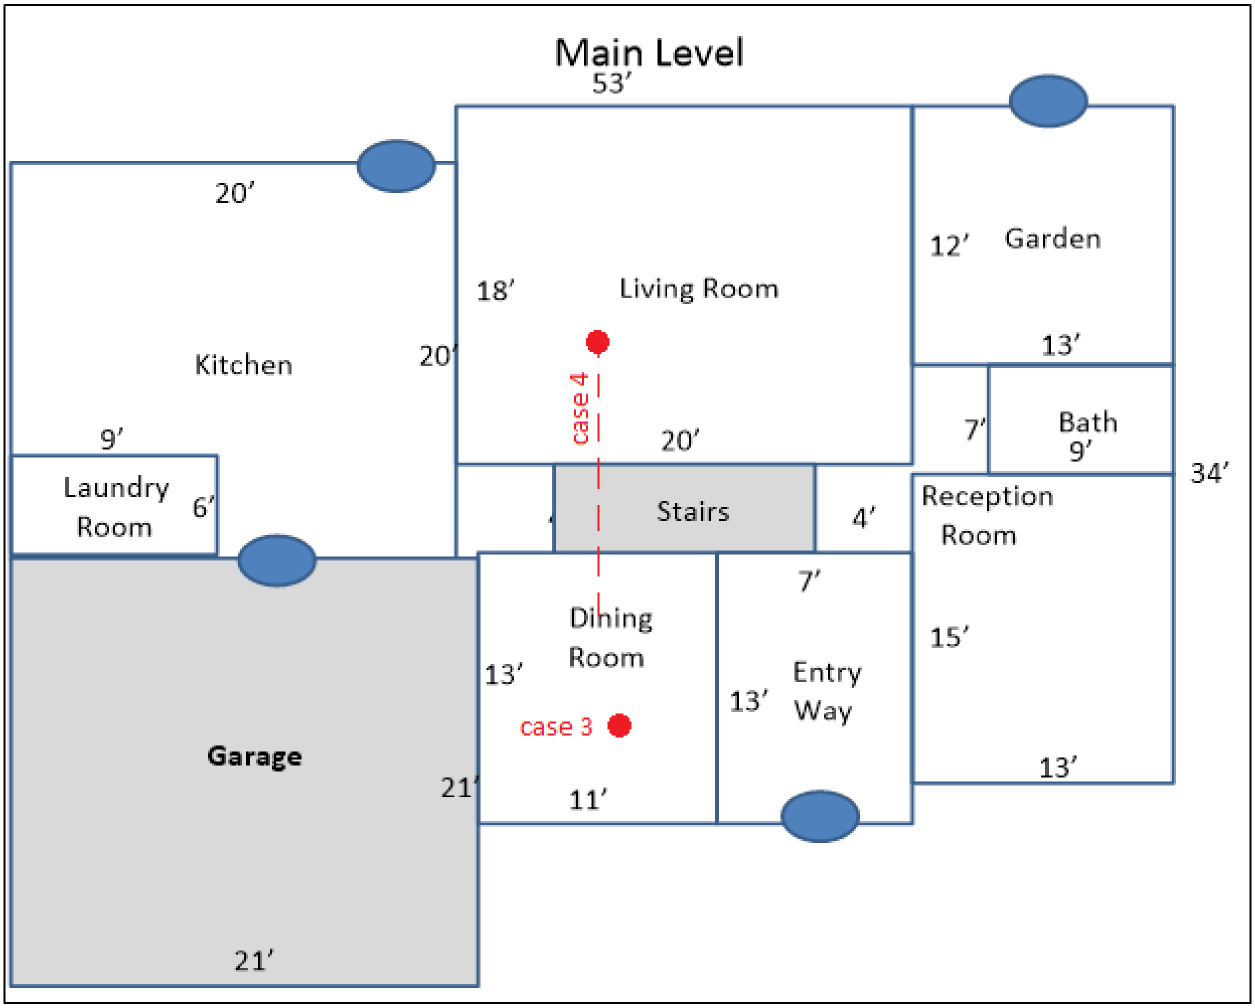

Figure 7 – WiFi test house, main/mid level floorplan with test cases

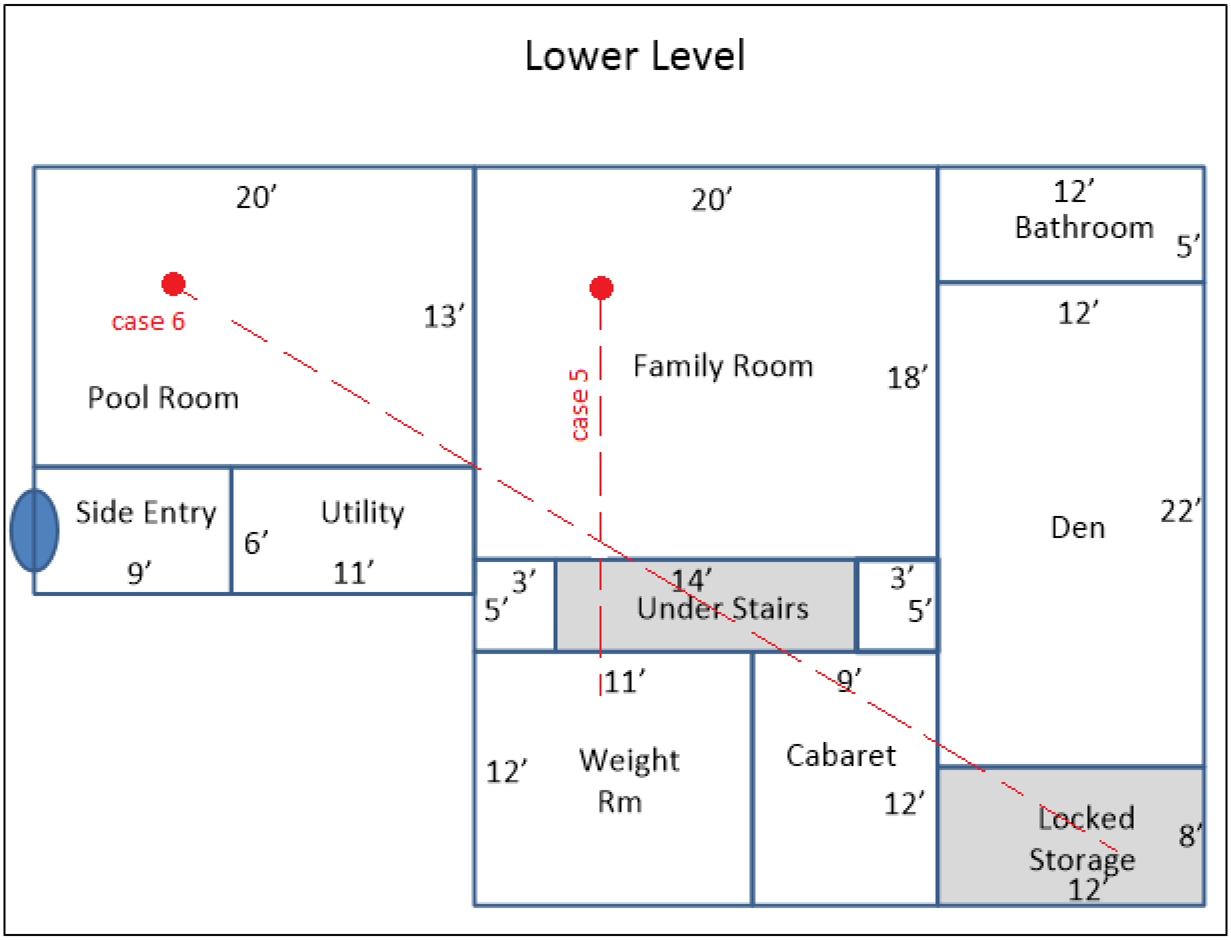

Figure 8 – WiFi test house, basement level floorplan with test cases

To validate expectations for the trunk's performance at low power, this average home's coverage map (as received bitrate performance) was measured across multiple links spanning near-room distances to several lumped element wall/flooring transitions at distances up to 60 feet. Six test cases, each at 100 mW, 200 mW and 400 mW total power and 80 MHz of bonded channel bandwidth were conducted. The link tests were performed with available WiFi5 endpoints set for channel 153 in the U-NII-3 band (so that through-air and lumped transition losses would mimic – albeit slightly optimistically -- conditions in the 6 GHz band). The test results using a standard iperf3 reference (scripted to include TCP messaging overhead) between two 4-chain devices yielded the following results (note that the rates were inclusive of distinctive device radio behavior around MCS selection and AGC setpoints):

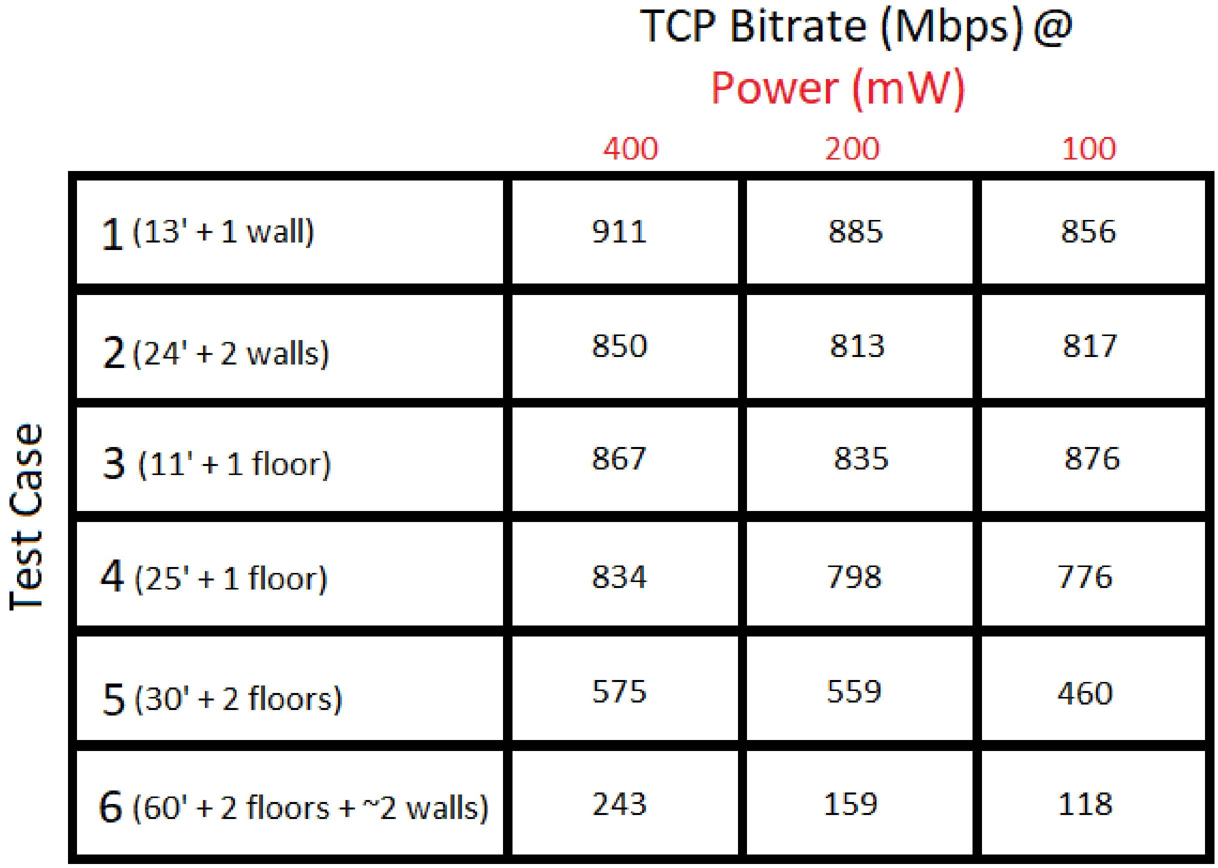

TCP Bitrate (Mbps) @ Power (mW)

Test Case	400	200	100
1 (13' + 1 wall)	911	885	856
2 (24' + 2 walls)	850	813	817
3 (11' + 1 floor)	867	835	876
4 (25' + 1 floor)	834	798	776
5 (30' + 2 floors)	575	559	460
6 (60' + 2 floors + ~2 walls)	243	159	118

Figure 9 – TCP bitrate performance across six test cases in WiFi house

One cautionary observation: the bitrate best-case asymptote of 911 Mbps, being TCP, can be viewed as just over 1 Gbps UDP (accounting for approximately 10% TCP signaling overhead). This is shy, however, of the expected 1.4 Gbps UDP rate expected at best MCS for 4SS WiFi 5 devices (refer to Figure 13) and appears to be an implementation artifact of the devices used in testing. (Such was confirmed by rescripting the channel BW for 40 MHz and observing a drop in bitrate by a factor of exactly 2.) But it is encouraging that 400 mW manifests the ability to light up a client device to over 200 Mbps service at 80 MHz of channel bonding (as pointed out, measured with nearly 30% implementation overhead and at TCP) across the longest (two-floor breaching) diagonal reach in the study.

Cases 2 and 4 suggest that the loss through flooring approximates that of 2 interior drywall transitions (1 floor ~= 2 walls). Regressing the measured data in case 5 against the analytical expectations in Figure 13 (below) to produce the floor transition loss goes as follows: 1) Free-air path loss of 5.765 GHz at 30 feet amounts to ~ 67 dB; 2) Figure 13 references loss at one meter (48dB) so move along the x-axis to 67-48 or 19 dB; 3) the reference curve assumes 1W of power but case 5/400 mW means we move an additional 4 dB to the right (to account for the lower power of the test) – so the operating point in free air would be here at 23 dB path loss; 4) the test case showed TCP performance at 575 Mbps with the hardware used so now we adjust for that implementation: multiply the TCP rate by 1.1 to get UDP, then by 1.4 to overcome the implementation loss in the HW used, so 575 x 1.1 x 1.4 = 886 Mbps; 5) Traverse the bitrate curve in Figure 13 down to the point where it indicates ~890 Mbps or so and drop down vertically to read the operational path loss (~41 dB); 6) the operational path loss minus the free air path loss indicates what the transition losses are, so 41-23 = 18 dB. Since this is accounted for by 2 floor transitions, this yields 9 dB for a floor loss. Floor loss being approximately twice the wall loss (in dB), this implies 4.5 dB for each drywall. Both of these estimates compare well with the extant literature on 2.4 GHz and 5 GHz indoor material transmission losses (typically in the 3-4 dB range for drywall and 6-10 dB across the bands from 2.4 out toward 6 GHz).

A USC "common material building loss" study (refer to reference section) from 2002 provides some additional frequency dependent data across the entire 2.4 – 7 GHz region. In broad summary, interior drywall transmission losses in that study averaged a fairly constant value over the bands in question and typical wooden beam and plywood flooring transmission loss seemed to monotonically increase from around 5 dB at 2.4 GHz to nearly 10 dB at 7 GHz.

To complement the measured data (and provide some calculus for expectations of bitrate performance that 6 GHz and similar power should be much better than these measured results), an analytical expectation for free air service radius of 4-chain WiFi6 endpoints operating in similar fashion as the test cases shows the following UDP bitrate expectations across increasing link losses at two power points:

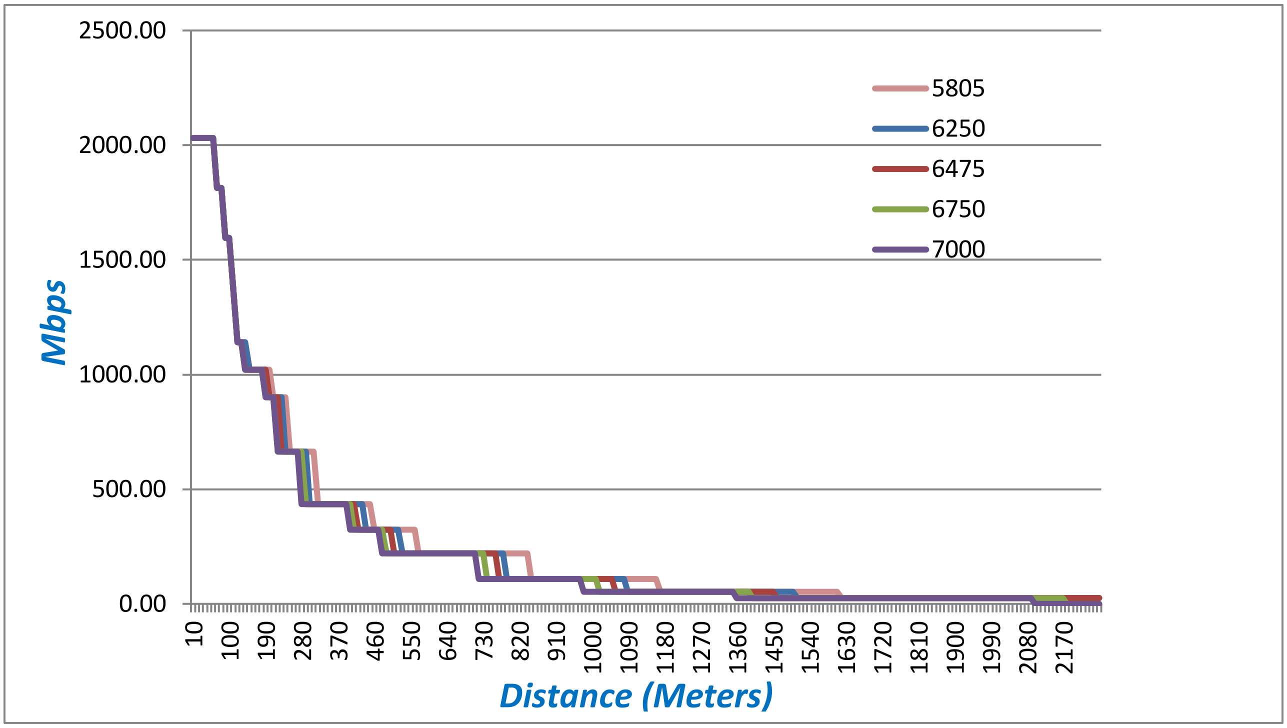

Figure 10 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band, 250 mW

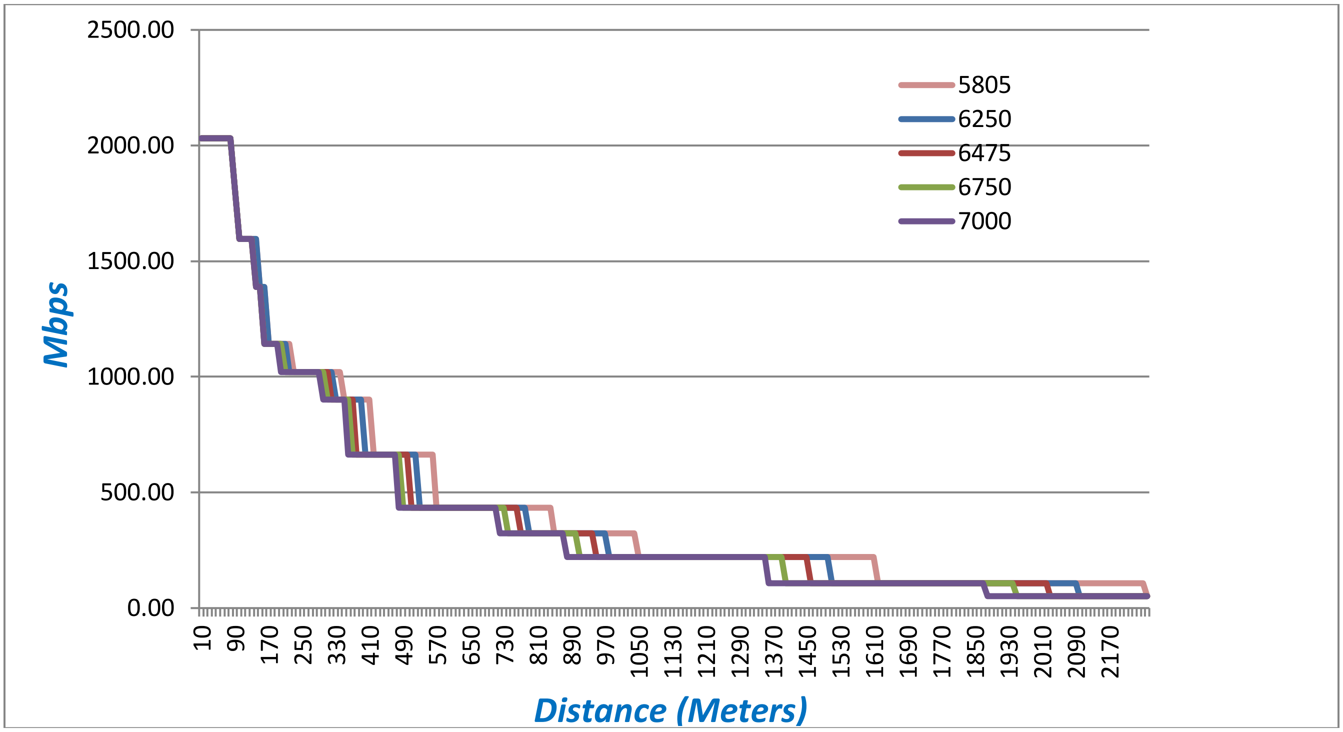

Figure 11 – 4 SS, 80 MHz UDP bitrate service radius at 5 frequencies in the 6 GHz band,1W

2.2. Why low power?

With legacy 802.11 indoor radiated power limits permitting more robust levels, it is reasonable to wonder why one might be interested at indoor performance achievable with a very modest 250 mW EIRP for the 6 GHz band. The answer lies in good neighbor coexistence (especially with extant outdoor 6 GHz infrastructure coupled with a desire to be conservatively biased with respect to interference) and the potential for the FCC to partition up the 6 GHz band into “standard-power” (presumably outdoor or outdoor/indoor uses) and “low-power” (indoor) bands in the first place. Of keen interest at the designated low power bands is the ability for indoor CPE to jettison intervention from cloud (or edge) -based interference arbitration schemes associated with the Automated Frequency Coordination (AFC) function (itself bearing some similarity to the CBRS band’s Spectrum Access System -- SAS). Refer to the NPRM proposal under consideration at the time of this writing:

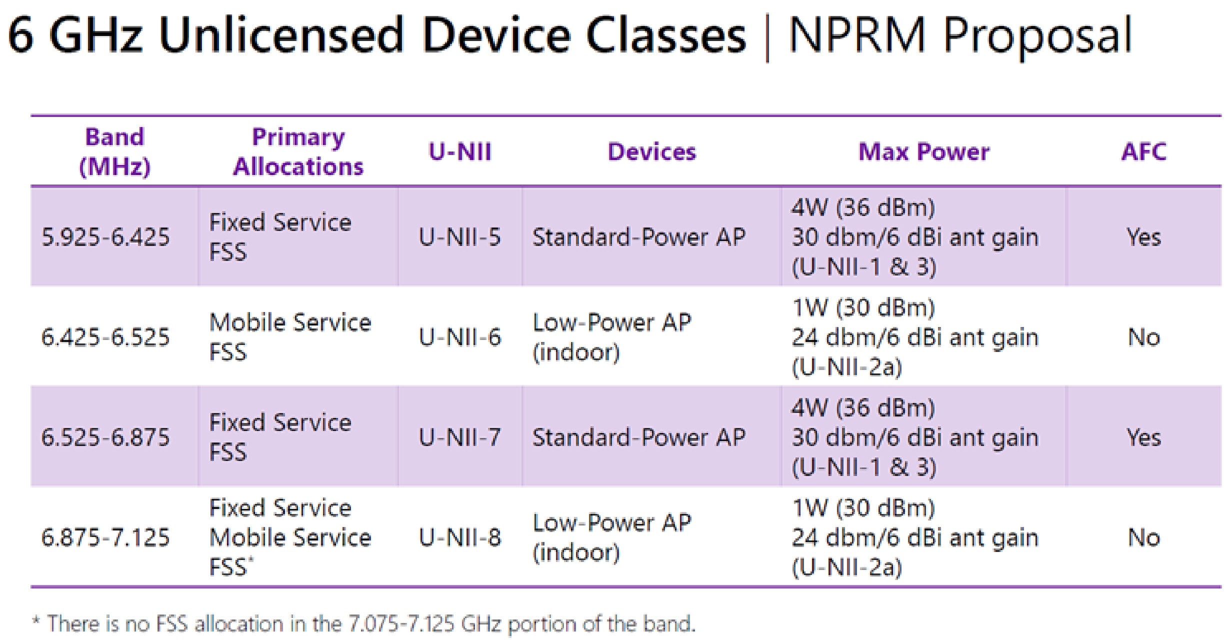

6 GHz Unlicensed Device Classes | NPRM Proposal

Band (MHz)	Primary Allocations	U-NII	Devices	Max Power	AFC
5.925-6.425	Fixed Service FSS	U-NII-5	Standard-Power AP	4W (36 dBm) 30 dbm/6 dBi ant gain (U-NII-1 & 3)	Yes
6.425-6.525	Mobile Service FSS	U-NII-6	Low-Power AP (indoor)	1W (30 dBm) 24 dbm/6 dBi ant gain (U-NII-2a)	No
6.525-6.875	Fixed Service FSS	U-NII-7	Standard-Power AP	4W (36 dBm) 30 dbm/6 dBi ant gain (U-NII-1 & 3)	Yes
6.875-7.125	Fixed Service Mobile Service FSS*	U-NII-8	Low-Power AP (indoor)	1W (30 dBm) 24 dbm/6 dBi ant gain (U-NII-2a)	No

* There is no FSS allocation in the 7.075-7.125 GHz portion of the band.

Figure 12 – Pending FCC NPRM showing consideration of non-AFC low power bands

If one could show impressive indoor service reach with power at (or better still, 6 dB below) FCC considerations for indoor use, it follows that the indoor application of 6 GHz to augment present-day 2.4 and 5 GHz bands should slot in with little concern for macro-scale microwave infrastructure interference (and this, without resort of WiFi6's impressive downstream – and with Wave 2 devices, upstream – spatial directivity represented by MU-MIMO). Switching consideration to MDU structures, the general rule of thumb that "the minimum necessary power required to sustain link throughput is the power level at which one should operate" does its part to minimize housing unit-to-unit interference potentials. (More about this in a subsequent section.) To illustrate WiFi6's inherent throughput advantages over WiFi5 (and perhaps dispel some entrenched legacy concerns regarding how much power is necessary), consider the following chart showing WiFi6 performance radius at 250mW versus WiFi5 at 1W:

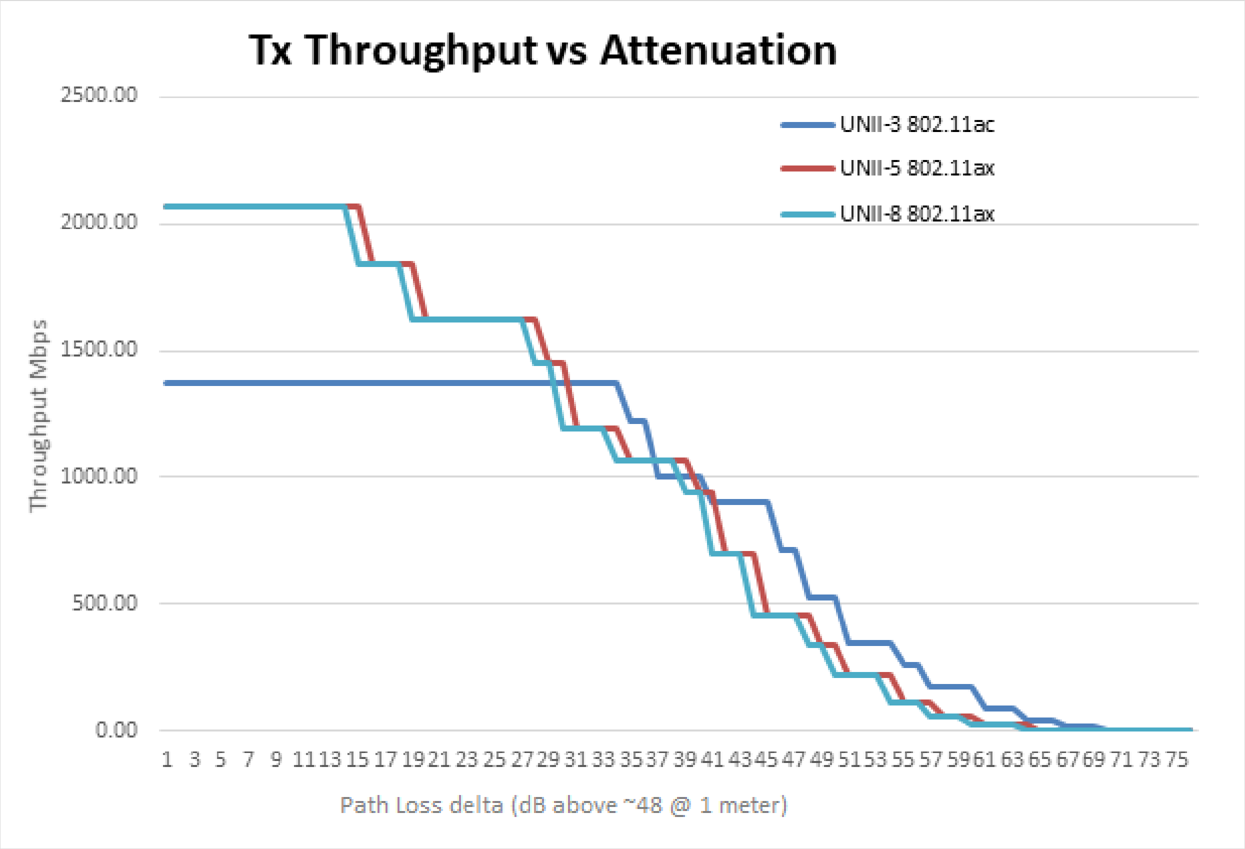

Figure 13 – 250 mW WiFi6 UDP bitrate curve @ U-NII-5 & -8 versus 1W WiFi5 @ U-NII-3

As can be seen, the available receiver sensitivities and ability to support denser spectral modulation schemes of the WiFi6 links imply that for expected indoor AP spacing and coverage, the 6 dB lower power WiFi6 links do as well – and much better, at close range -- as the higher power WiFi5 reference link. The crossover point shown in the chart occurs at nearly 1.4 Gbps UDP service and at a total path loss of 77 dB or so (representing a throw of 85 feet, absent any floor or wall transition losses). So clearly, indoor RF considerations for WiFi6 need not be inclusive of 1W power levels – which, aside from the interference concerns previously noted, bodes well for HW implementation considerations for overall device dissipation (and ultimately, cost).

2.3. Challenging the WAN capacity

A high-bitrate demand scenario with mixed-epoch WiFi clients was crafted to illustrate the raw new capacity represented by only partial exploit of the 6 GHz band. In this exercise, a WAN attachment GW device is wirelessly trunked with a quad-band (2.4, 5Lo, 5Hi, 6 GHz) extender to examine traffic capacity of the trunked link. For purposes of the study, the 6 GHz trunk is deemed to be supported by a 4 x 4 radio scheme at that band (though exploits up to 8 chains are permitted). The near-field meshes associated with the WAN GW are not considered (the assumption being that these data demands are supplied and scavenged directly at the GW/WAN aperture and would not impact performance considerations for the trunked extender and its separately served mesh of clients). The extender is linked to the GW via a 160-MHz wide, bonded 6 GHz channel of "best effort" WiFi priority and this AP sees a mix of traffic it manages with five end devices (three specific 6 GHz/WiFi 6 clients with defined spatial capacity and two "ensemble" devices, representing proxy traffic to multiple 2.4 and 5 GHz clients at the specified bitrates and link priorities). Three new WiFi6 / 6 GHz clients are shown in the accompanying figure as links L1, L2 and L3. L1 is a connection to a UHD TV which travels 50 feet through 1 floor and 3 wall transitions and exploits a 2 SS radio with 40 MHz of bonded channels to deliver ultra-high resolution and definition video via a 70 Mbps stream. (This is a massive ask for a streaming client.) L2 represents a 2 SS mobile device connection 20 feet from the extender, through one wall. And L3 proxies an HD video stream to a TV or tablet 3 rooms (40 feet, 3 walls) away from the extender supporting a 30 Mbps stream via 2 SS and 80 MHz worth of bonded channels. Link CL1 is a WiFi 5 multiclient proxy which, in ensemble, represents a 400 Mbps demand with VI priority located 30 feet away with 1 wall and 1 floor to transition and commands a 4 SS connection of 80 MHz BW. Finally, CL2 is a single spatial stream of BK priority with the same topographical impediments as CL1 but asks for 50 Mbps support on a single 802.11n, 2.4 GHz WiFi channel. This schema is representative of an existing WiFi-invested single-family detached home which is adding WiFi6 at 6 GHz as a service(s) expansion. A block diagram of the exercise follows:

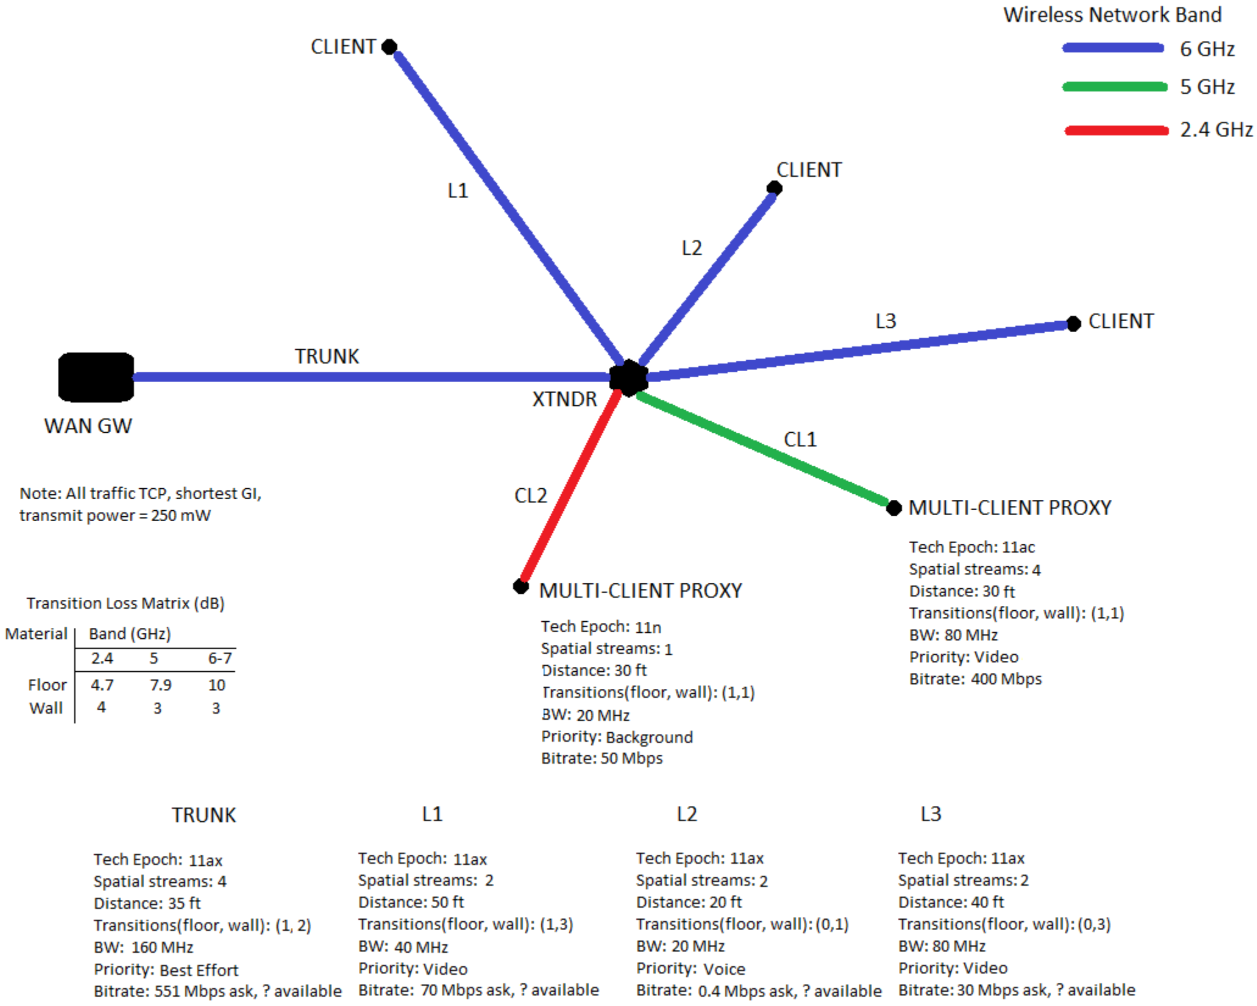

Note: All traffic TCP, shortest GI, transmit power = 250 mW

Transition Loss Matrix (dB)

Material	Band (GHz)		
	2.4	5	6-7
Floor	4.7	7.9	10
Wall	4	3	3

MULTI-CLIENT PROXY
Tech Epoch: 11ac
Spatial streams: 4
Distance: 30 ft
Transitions(floor, wall): (1,1)
BW: 80 MHz
Priority: Video
Bitrate: 400 Mbps

MULTI-CLIENT PROXY
Tech Epoch: 11n
Spatial streams: 1
Distance: 30 ft
Transitions(floor, wall): (1,1)
BW: 20 MHz
Priority: Background
Bitrate: 50 Mbps

TRUNK
Tech Epoch: 11ax
Spatial streams: 4
Distance: 35 ft
Transitions(floor, wall): (1, 2)
BW: 160 MHz
Priority: Best Effort
Bitrate: 551 Mbps ask, ? available

L1
Tech Epoch: 11ax
Spatial streams: 2
Distance: 50 ft
Transitions(floor, wall): (1,3)
BW: 40 MHz
Priority: Video
Bitrate: 70 Mbps ask, ? available

L2
Tech Epoch: 11ax
Spatial streams: 2
Distance: 20 ft
Transitions(floor, wall): (0,1)
BW: 20 MHz
Priority: Voice
Bitrate: 0.4 Mbps ask, ? available

L3
Tech Epoch: 11ax
Spatial streams: 2
Distance: 40 ft
Transitions(floor, wall): (0,3)
BW: 80 MHz
Priority: Video
Bitrate: 30 Mbps ask, ? available

Figure 14 – Model of 6 GHz / 160 MHz BW in-home backbone trunk and service mesh

In addition to the bitrate demands, link losses due to distance and wall/floor transitions are included and link priorities established. The 802.11e priority levels (from highest to lowest) are voice (VO), video (VI), best effort (BE) and background (BK). The resulting whole-home wireless network performance is captured in this summary:

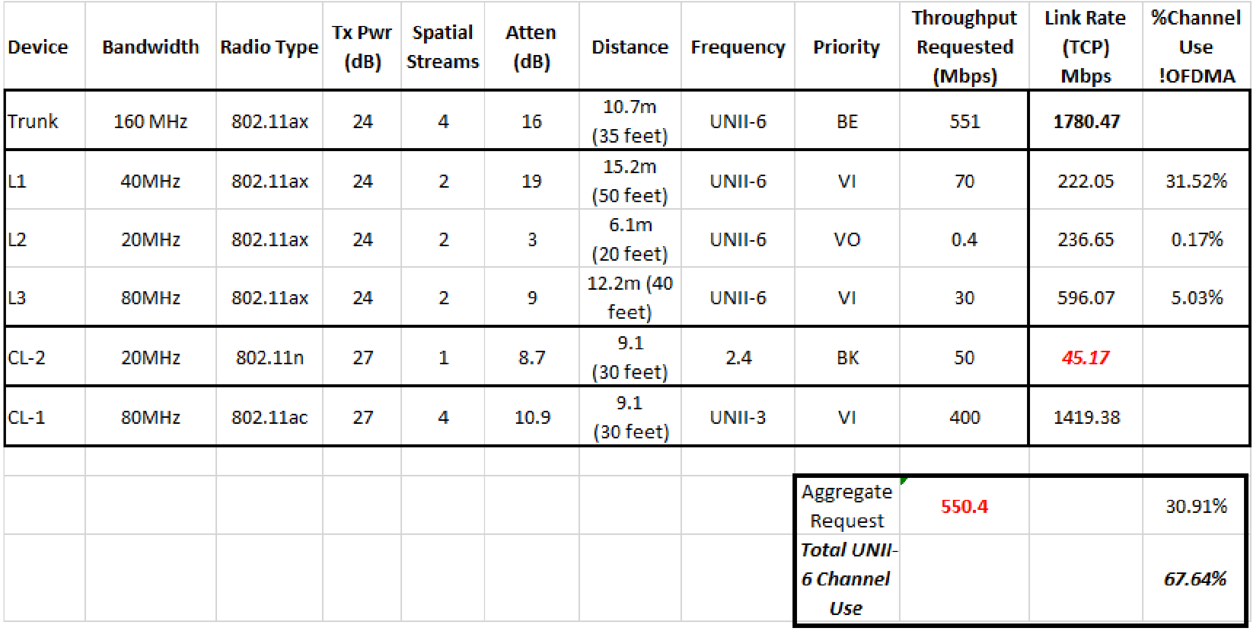

Device	Bandwidth	Radio Type	Tx Pwr (dB)	Spatial Streams	Atten (dB)	Distance	Frequency	Priority	Throughput Requested (Mbps)	Link Rate (TCP) Mbps	%Channel Use !OFDMA
Trunk	160 MHz	802.11ax	24	4	16	10.7m (35 feet)	UNII-6	BE	551	**1780.47**	
L1	40MHz	802.11ax	24	2	19	15.2m (50 feet)	UNII-6	VI	70	222.05	31.52%
L2	20MHz	802.11ax	24	2	3	6.1m (20 feet)	UNII-6	VO	0.4	236.65	0.17%
L3	80MHz	802.11ax	24	2	9	12.2m (40 feet)	UNII-6	VI	30	596.07	5.03%
CL-2	20MHz	802.11n	27	1	8.7	9.1 (30 feet)	2.4	BK	50	***45.17***	
CL-1	80MHz	802.11ac	27	4	10.9	9.1 (30 feet)	UNII-3	VI	400	1419.38	
								Aggregate Request	550.4		30.91%
								Total UNII-6 Channel Use			***67.64%***

Figure 15 – Expected Performance of the exercise model

Some significant aspects are immediately apparent: 1) though only fronthaul demand (~550 Mbps) is calculated, if symmetric duplex demand were placed on the trunk (though this is not a likely requirement with the video streaming use cases cited), the total available capacity at the link losses specified would amount to one and a half times the fully symmetric demand (1.1 Gbps) the network would then require; 2) all 6 GHz band traffic (trunk or any of the mesh links) can be distributed without contention (and with surplus BW available); 3) 6 GHz channel capacity on the employed bonding schemes is well below thresholds at which prioritized scheduling need occur (i.e., there are no queuing latencies aside from framing alignment which occur with any of the services mounted – data is dispatched as soon as it is received). Furthermore, as the trunk is a P2P link with management set by the WAN GW, the multi-client (OFDMA) benefits of WiFi6 are not a consideration for this haul – the GW's sole discretionary responsibility is to determine how much bandwidth it needs and where to locate the bonded channels. A note from the legacy support side on this simulation is that the single SS associated with the 2.4 GHz band and the requirement for 50 Mbps service there is that such cannot be supported (at 45 Mbps). This could be remediated by support of a second SS, more BW or an uptick in priorities – but it underscores the motivation for moving away from '11n and onto the MAC support offered by WiFi6.

Digesting this data and simultaneously acknowledging that incorporation of 6 GHz clients into home meshes will consume some fair amount of time, the key inference is that for single-family detached homes, the overwhelming availability of largely interference-free channel BW posed by the 6 GHz band and first scheduling resort of merely distributing trunked traffic in FDM fashion is an obvious source of significant WiFi performance improvement for the home. This further implies that consideration of queuing strategies for opportunistic packaging of the RUs (perhaps against a weighted judgment of total throughput and accumulated latency on a per-service, per-client basis) – versus merely dispatching data chunks immediately on the assigned channel -- will not be obviously necessary for several years.

MDUs, however, could represent a much different story – bearing in mind that, as pointed out above, initial deployments of band-compatible devices are likely to enter service in scattershot fashion, which will tend to minimize initial CCI emergence. The eventual, thirty thousand foot "qualitative analysis" view is that -- coupled with the potential for even low-power 6 GHz coverage to bleed beyond the intended coverage area of an access point -- MDU's forced spatial concentration of 6 GHz-compliant APs should test MU-MIMO, OFDMA and BSS coloring attributes of WiFi6 before these ever become necessary alternatives for detached single family dwellings.

And it is worthwhile to note that the 6 GHz spectral pool being manipulated by the APs in an MDU is a necessarily compromised resource. All of the APs compete for leverage of this asset by their clients based upon perceived availability (by either end component of the intended links) of useful spectrum. BSS coloring promotes re-use of spectrum as may be possible, but its very nature also guarantees that the spectrum to be used may be compromised to some degree. This is because with the shifted "channel in use" thresholds associated with discernment of competing traffic from another BSS color, there will inevitably be loss of fade (noise) margin due to ingressing signals on those channels the local AP deems available for use by its mesh. Channels then chosen with nonzero levels of ingressing energy would be forced to reduce spectral efficiency (due to modulation backoff to cope with degrading SINR) and this would necessarily drive down link bitrates (hence, reducing opportunities to pump data, potentially resulting in buffer growth on one end of a given link or the other). A reasonable question might be "How soon does this happen and how dramatic are the effects?"

3. Tinkering with expected performance in MDU environments

3.1. Napkin musings on the scope of the ask

Granting that the MDU scenario will present the most obvious residential indoor challenge to unlicensed exploit of the 6 GHz band, it seems prudent to attempt to put some numbers to help define the magnitude of any problematic issues. Though not without significant deviation in the data, the "average" MDU in the US comprises 12 units in a two- to three-floor building, each unit of roughly 900 square feet – and most of these one or two bedroom. Floorplans obviously vary, but we might assume 36' x 25' units with a centrally located AP and a maximum of 3 wall transitions to service each unit's palette of 6 GHz clients. To drive AP concentrations up, we can go to 150% of average and conjure an 18-unit MDU with 3 floors of 6 units each.

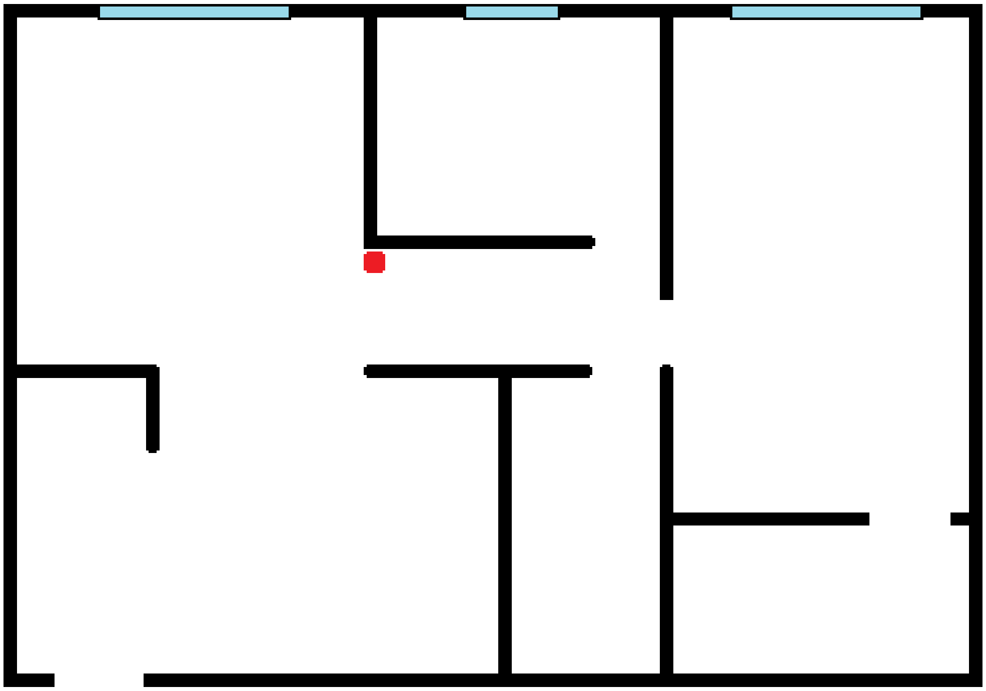

Figure 16 – Floor plan of "typical" 900 square foot apartment, AP location in red

Regarding the units, we might project 2 occupants, each with 3 simultaneously-used client devices (mobile, tablet or PC and a TV) so there are 6 active client devices per unit. To continue to drive the worst case, we can perhaps assume ~240 Mbps per unit (2 x 70 Mbps UHD video streams, 2 x 30 Mbps HD video streams, 2 x 20 Mbps browsing).

3.2. Crunching some numbers

The overwhelming concern in MDUs is the potential for CCI from neighboring units to aggregate towards the middle of the building. A single floor is shown below. The full impact of near neighbor energy would occur in the middle unit(s) of the middle floor, where ensemble distances to either end (AP or client) of all other wireless links in the MDU would be shortest (and hence, CCI strongest):

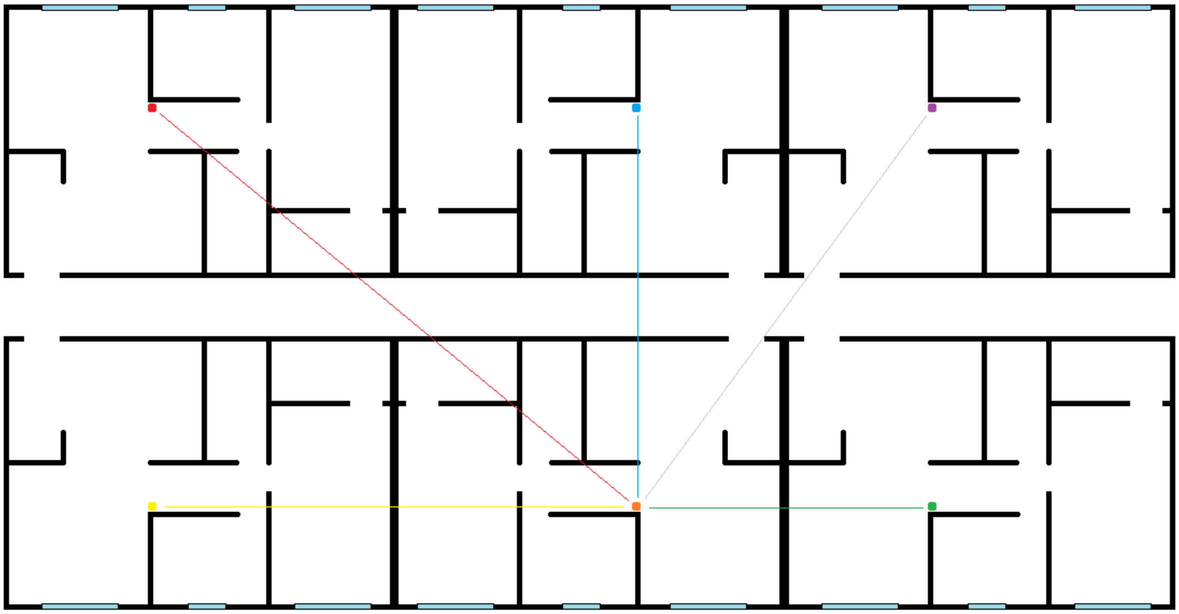

Figure 17 – One floor of example MDU (6 units) showing CCI peak location

With 18 units demanding 4.32 Gbps in ensemble (assuming everthing is unicast), the issue becomes how to route the demand with 400 MHz of low power WiFi BW at U-NII-6 and U-NII-8 (10.8 bps/Hz implied). The employ of multiple SS immediately rescues this; in fact using only half of the 8-chain WiFi 6 specification on spatial diversity (4 SS, then), one could theoretically quarter the stated spectral density demand to a very manageable 2.7 bps/Hz (and still not be considering any OFDMA exercise). Put another way: if each (4 SS) AP is considered atomically, 240 Mbps is child's play. Referring to our assumptions on unit topology, a worst-case service radius inside the unit might be 22 feet with 9 dB worth of drywall transition losses. The path loss represented would be 75 dB (66 dB through-air and 9 dB lumped). Recall from Figure 13 that a 250 mW, 4SS WiFi 6 AP with 80 MHz of channel bonding would produce roughly 1.6 Gbps UDP throughput at 75 dB path loss. So even if this device restricts its spectral use to a single 20 MHz channel, 1600 / 4 = 400 Mbps (167% of stated demand) should be available to its clients throughout the unit floorplan.

Now, if each AP only requires a single 20 MHz channel to meet its requirements, then there is enough U-NII low-power BW so that each AP in the MDU could exercise its own 20 MHz channel (18 x 20 = 360 MHz used) without contention for spectrum. One may complain that this is merely fortuitous or somehow a product of a prescient (or perhaps historically aware?) cloud provisioning agent privy to the link data from all of the APs. After all, such peak demand would be the result of arbitrarily originated and continued sessions – which would bloom in aggregate and decay over the course of a day. And there is nothing in the 802.11 specification which restricts an AP to anything other than random pursuit of available channels. But this result could also be an end product of APs whose behavior was such that they routinely mined currently used channel BW for as much client data transport as required before venturing farther afield (spectrally speaking) and were rigorous about initiating connections with the minimum bonding necessary.

There is also a management opportunity here for even more power backoff in the MDU case to improve coexistence implications (within and outside the structure) and still have bursty data overhead in the running links. After all, the service throws per unit are much more constrained than for the detached single family dwelling and there are no floor transitions – and fewer walls – involved. For example in the modeled MDU, the service path loss is estimated at 75 dB and the goal for the unit is 240 Mbps. Referring again to Figure 13, if we drop the power to 100 mW and maintain the same 4 SS AP with 80 MHz channel bonding, we essentially shift our operating point 4 dB to the right; the rate is 1200 Mbps. At a single 20 MHz channel, this still produces ~ 1200 / 4 or 300 Mbps. In fairness, in the game of "WiFi power chicken", such backoff would need to be practiced by all wireless partners to be most effective, but the end results are significant.

So if available channelization suffices to supply all wireless device connectivity without spectral overlap, as in this particular MDU example, it follows that there are no obvious one-hop latency contributors (due to access contention or CCI, say) past two-way dispatch and recovery mechanisms at each of the two link endpoints; so no backoffs are required. Such latency ought to amount to less than 200 usec then, (roundtrip) for the MAC priorities VO, VI or BE. Relative to WiFi legacy performance in the crowded 2.4 or 5 GHz bands, this pins potential latency in the 6 GHz band to the "extremely responsive" bin.

3.3. The Power of BSS Coloring in MDUs

Fortuitous or no, this crafted MDU exercise did not test the most noteworthy of WiFi 6's MAC benefits, only the bounty which is 6 GHz unlicensed BW. It appears near-future 6 GHz client device populations in indoor residential environments do not look to impose sufficient contention to exercise these MAC mechanisms -- admittedly designed primarily to service large-venue WiFi scenarios. But we can perhaps calculate more Draconian (if unrealistic) indoor plays which would challenge the WiFi 6 MAC if we are willing to significantly up the service loads (as ensemble bitrate and number of separately served clients) past what one would "normally" expect for residences.

If one examines the proposed low power portion of the 6 GHz band for maximal unlicensed bitrate carriage using WiFi 6 mechanics, the asymptotic numbers are impressive: the combined BW of U-NII-6 and U-NII-8 could yield nearly 24 Gbps and 1,440 1 x 1 client devices per BSS served by an 8 SS AP (delivering over 15 Mbps service continuously to each of those clients at close-in service radii). In practice, of course, such is nowhere near achievable; SINR pollution from the OBSS populations and microwave rogue ingresses, differential service radii from AP to clients, non-isochronous session behavior, etc etc all conspire to compromise delivery efficiencies of the MAC and PHY. And on the other side of the coin, there are any number of mobile applications (phone calls representing a classic example) which require much less than 15 Mbps data connectivity -- which WiFi 6's OFDMA support could interleave into additional client support (over 60,000 simultaneous voice calls, say).

In qualitative terms, the onset of contention issues are guaranteed to occur when simple FDM'ing of the available BW cannot be assured and the bitrate ask on the served channels exceeds asymptotic capacity. But this represents a massive concentration of disparate BSS domains and clients (the former more than the latter, given the implications of lack of control of the OBSS clients). The LPI BW portion of the 6 GHz band being considered at this point in time is 400 MHz. Restricted to a (lucky) distribution of 20 MHz channels, this implies once more than 20 AP's on different BSS colors are operating such that they (or their client populations) are within ~10 dB service radius of all abutting OBSS domains, various backoffs can begin to occur. This effect is accelerated the fewer spatial streams employed by the APs in question and the higher the bitrate (BW) demand – but the onset is anything but precipitous. Reference

Figure 13 again at the 31 dB listed operating point (75 dB path loss @ 100 mW) for an "average" MDU unit and notice that even though MCS backoff is noticeable over the next 10 dB of SINR degradation, the rate is no worse than about half what was sustained before. Unless this attenuation in rate pushes the AP to search for additional BW, the service would simply carry on.

And this type of contention presumes that near-neighbors would not associate interference with close-abutting alternate BSS domains in the first place and fail to seek out better spectral alternatives based upon appropriately shifted channel-in-use detection (they would – that's the very point of detecting color collision and reassigning colors). That's the mechanism of OBSS discerned channel re-use – and the threshold of this channel re-use can be shifted to expose the level of interference the AP is willing to entertain.

Given the minimal physical dimensions of the service spaces (call it ~7200 cu ft) and the cubic presentations of the sources of CCI in the 3D model of MDUs, one might suppose that the seven nearest 3D neighbors would pose the greatest threats. (Immediately above and below the unit in question and the five same-floor surrounding units). Note that other units may interfere, but the CCI implications begin to fall off dramatically with increased distance – and surface transitions – that accumulate). If all bitrate consumption in the MDU is high and evenly distributed, then an 8SS AP could still be bonding just a fraction of the available 400 MHz of LPI and, even with huge backoff on MCS (to 10-15 dB of impairment), be capable of delivering over 1 Gbps to its clients due to appropriate repurposing of the "colored" channels available to it.

3.4. External Interference by Inside WiFi at 6 GHz

It is worthwhile to close with a couple of observations on WiFi radiation which escapes residential low power deployments and "enters the wild" (inside/out propagation) since there are some number of legacy outdoor installations which utilize the band (as P2P or P2MP links) and new sources of interference for their operation would not be welcomed with much enthusiasm. To be conservative is estimating impact, one can neglect the selectivity of the link antennas used there (although certainly, the patterns employed are usually very much less than even hemispherical in nature and engaging the antennas on a back pattern would result in ample relative attenuation – front-to-back ratios easily exceeding 40 dB+ in most applications).

Indoor WiFi leakage radiation at 6 GHz needs to overcome the service throw indoors as well as the ultimate transition through exterior walls. In examining outside-in behavior for CBRS, transition loss for wood paneling was pegged at around 6 dB, brick or HardiPlank at 13 dB and stone at around 25 dB. In this case we will reference the loss of wood siding in the calculations, to maintain worst-case parameters for consideration. And based on our previous discussion of detached residences versus MDUs, it appears that the latter can operate with lower WiFi service power with a WAN attachment near the middle of the unit and the former might feature a WAN GW element just inside the exterior wall, operating at higher power. So for the calculations, we will reference 250 mW of WiFi source power with only a 6 dB offset to cite for knockdown of interference radius. Furthermore, we will assume no significant vegetation nor topographical impediments around the candidate interferer.

As to the outdoor equipment, we will assume that the receiver/antenna combination used operates at a convenient 20 MHz BW – but inconveniently on the same frequency as the WiFi in question (and this, at the very bottom of the 6 GHz band), has a 6 dB NF receiver and anticipates worst-case operation down to 8 dB SINR at the detector. (Lots of presumptions here, admittedly). These restrictions set the edge of the operational fade budget for the microwave tower down to a signal level of around -87 dBm. Any interference would need to be at least 10 dB lower (more preferable) which suggests our leakage radiation needs to be at or below -97 dBm at the antenna to be confirmably non-invasive.

If we hold to our chain of worst-case presumptions, this puts any microwave tower within ~ 2.3 km of the home in question at risk for interference. For perspective, however, if the WiFi AP were operating at 100 mW in the middle of a stone home, itself in a mildly forested area, this radius would collapse to 30 meters (mostly due to internal wall transitions, the huge signal cost of exiting a stone exterior home and tree/vegetation impedances which can easily reach 10 dB or more). A geometric mean of 260 meters is the result of these extreme considerations (though again, this is without allowance for antenna selectivity).

To frame this with radiation patterns in mind, one could assume a tower microwave with a 120 degree main lobe and rationalize the home as having perhaps a hemispherical pattern (referring to the placement of the AP against the exterior wall, the differential internal transitions at minimum would describe a non-isometric pattern of some sort with perhaps 12 dB or so of minimal directive imbalance). If the azimuth offsets of the orientations of the tower and the home were randomly distributed, this implies that perhaps one time out of six when distance to a tower from a home was within the interference radius, the relative directivities of the two radiation patterns would conspire to create CCI.

In generic terms, one would perhaps characterize this as a very slight chance of interference. (And to be sure, the dual of the situation holds – what passes for CCI at the microwave receiver would likely be seen as a “busy channel” by the WiFi – which would eschew – or downrate -- its selection for use in the first place.)

Conclusion

The 6 GHz band provides a magnificent aperture for unlicensed wireless services to grow in a disciplined and future-proofed direction, providing instant relief for the capacity exhaustion and contention-based loss of efficiency occurring in the legacy 2.4 and 5 GHz bands while also delivering a self-installable, low-latency remedy which grants consumers an indoor-propagation-friendly wireless framework with massive, extensible bitrate support. In detached single-family dwellings, just the insertion of a dedicated 6 GHz / WiFi6 wireless trunk from WAN GW to middle-of-the-home extender promises multi-Gbps, interference-free whole-home WiFi coverage. MDU applications at 6 GHz, while more challenging due to the client and BSS densities implied by multiple overlaid networks operating in close proximity, nonetheless will greatly benefit from the 6 GHz spectrum access while barely invoking MU mechanisms in antenna directivity and the two-axes options for transmission data packing afforded by OFDMA. And the rather more friendly confines of single (one-floor) units in these structures mean additional low power backoff (perhaps to 100 mW, or less) can be implemented without concern for compromised services delivery.

Abbreviations

AP	access point
AFC	Automated frequency coordination
AGC	Automatic gain control
BE	Best effort
BK	Background

bps	bits per second
BSS	Basic service set
BW	Bandwidth
CBRS	Citizens broadband radio service
CCI	Co-channel interference
CPE	Customer premise equipment
dB	decibel
E2E	End-to-end
EIRP	Equivalent isotropically radiated power
FCC	Federal communications commission
FEC	forward error correction
Gbps	Gigabits per second
GHz	Gigahertz
GW	Gateway
HFC	hybrid fiber-coax
HD	high definition
Hz	hertz
ISBE	International Society of Broadband Experts
ISM	Industrial, scientific and medical
LPI	Low power implementation
MAC	Medium access control
Mbps	Megabits per second
MCS	Modulation and coding scheme
MDU	Multiple dwelling unit
MHz	Megahertz
MIMO	Multiple-in and multiple-out
MU	Multiple User
mW	Milliwatt
NF	Noise Figure
NFC	Near field communications
NPRM	Notice of proposed rulemaking
OFDMA	Orthogonal frequency division multiple access
OOB	Out-of-band
RF	Radio frequency
SCTE	Society of Cable Telecommunications Engineers
SINR	Signal-to-noise-plus-interference ratio
SNR	Signal-to-noise ratio
SS	Spatial Stream
TCP	Transmission control protocol
Tx	Transmit (or transmission)
UDP	User datagram protocol
U-NII	Unlicensed National information infrastructure
VI	Video
VO	Voice
W	Watt
WAN	Wide area network

Bibliography & References

IEEE P802.11ax/D4.2: *Draft Standard for Information technology – Telecommunications and information exchange between systems Local and metropolitan area networks – Specific requirements;* LAN/MAN Standards Committee of the IEEE Computer Society Copyright 2019

SCTE-ISBE and NCTA: *The New Home as a Hotspot: Wi-Fi Meet CBRS LTE and Meet Your Long Range Brother LoRA;* J.R. Flesch et al, Copyright 2018 SCTE-ISBE and NCTA

Propagation Losses Through Common Building Materials 2.4 GHz vs 5 GHz; Robert Wilson, USC 2002