Why 6 GHz Standard Power Wi-Fi is the Game Changer for Residential Use in the US

A Technical Paper prepared for SCTE by

J.R. Flesch
Director, Advanced Technology, Home Networks
Commscope
3871 Lakefield Drive, Suwanee, GA 30024
jr.flesch@commscope.com

Charles Cheevers
CTO, Home Networks
Commscope
3871 Lakefield Drive, Suwanee, GA 30024
Charles.cheevers@commscope.com

Kurt Lumbatis, Commscope

Bryan Pavlich, Commscope

Table of Contents

Title	Page Number

List of Figures

Title	Page Number

1. Introduction

We have an exemplary beachhead for better in-home Wi-Fi coverage with the instantiation of a low power indoor (LPI) effective isotropic radiating power (EIRP) spec (5 dBm/MHz power spectral density (PSD)) for access points (APs) by the FCC which promises a Gbps+ link budget for in-home data distribution (and scavenging) over 6E wireless networks. But the FCC also opened the door to standard power in that band (up to 36 dBm EIRP) for indoor use with the adoption of a spectrum coexistence scheme it refers to as automated frequency coordination (AFC). Using AFC and the higher power, we can now consider trunked indoor links which meet (or better) a 2.5 Gbps PHY for even large floorplan homes and perhaps enable mixed-power and mixed-band in-home mesh architectures. This enthusiasm is tempered somewhat by the observation that battery-powered clients are currently restricted by silicon to 20 dBm footprints at 6 GHz and as such, may determine extender density in the home for particular service mounts. However, alternating current (AC)-powered clients capable of power upticks (above LPI even if shy of standard power) – which include set top boxes (STBs), video streamers and gaming hardware – imply that we should be able to exploit link modulation and coding scheme (MCS) to fairly dense quadrature amplitude modulation (QAM) spectrum efficiencies for streaming and gaming services to these fixed endpoint clients – especially those of a heavy downstream data delivery bias.

This paper will examine the opportunities to be found in a standard power Wi-Fi 6E regime, discuss AFC implications (for cloud portal and AP endpoints) and suggest possible in-home architectural leverages of this substantial uptick in AP EIRP.

2. The LPI Reference Point

2.1 Availability of FCC-compliant Test Devices

6E consumer premise equipment (CPE) has begun to enter the market and it is obviously relevant to gauge the impact of LPI's allowable power in real scenarios which feature appropriate spatial stream (SS) scaling alongside the 6 dB client "power penalty" between client devices and APs -- and against path losses normally associated with various endpoint locations in a common floorplan. As in previous years, the 5300 square foot Commscope/Arris Wi-Fi house affords us the opportunity to test wireless links which emulate endpoint placements in homes up to that footprint in size and explore rate/reach and latency performance against these placements.

An available 6E-compliant smartphone was tested against a multi-band 4x4 AP to evaluate rate/reach throughout the Wi-Fi house. The results of these tests are shown below.

2.2 Early 6E Device Performance Data

Measurements were conducted early last year for nascent 6E devices (4x4 AP and 2x2 smartphone). These yielded the following measurements in the house at close operating range:

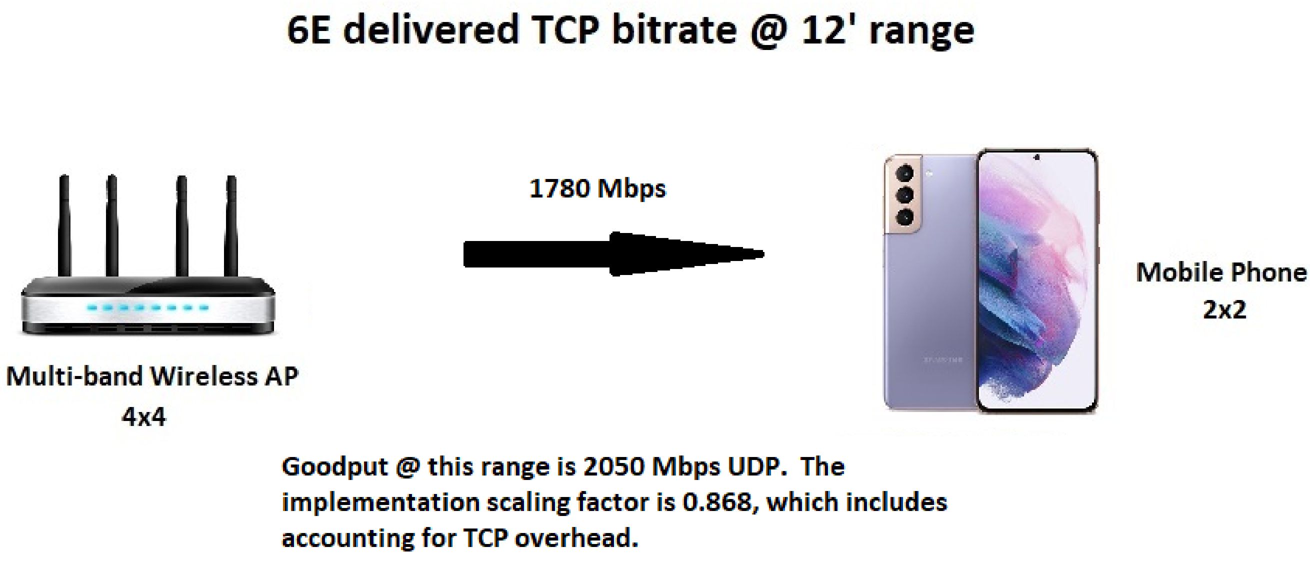

Figure 1 - Measured Close Range 6E LPI TCP Performance Between the AP and a Mobile Client

The testing was conducted across an open room at a radius of 12 feet or less (to establish a low path loss but far-field result which was expected to produce maximum MCS). The mobile was operating at backed off EIRP, relative to the AP, as mandated by the FCC. Free-space path loss (FSPL) at mid-band 6 GHz and that distance amounts to 60 dB (~49 of which gets you 3 feet off the antennae). Referring to the bitrate waterfall curve (goodput rate versus path loss) for 4x4 <-> 2x2 (albeit at equivalent powers, AP and client), we see that the link should support a goodput (bitrate accounting for Wi-Fi framing) of just over 2000 Mbps, with a TCP accounting moderating that to around 1800 Mbps or so (using 10% TCP overhead as a reference). Note that battery-powered mobile clients are likely limited even more -- to a maximum of 20 dBm EIRP -- which pushes the connection penalty to at least -7 dB relative to the AP. The operating points are captured on the waterfall curve for Wi-Fi 6 bitrates at LPI for 2x2 clients, leveraging the maximum 160 MHz channel bandwidth (note that the abscissa of the graph shows path loss off the antenna – which means you are actually down 49 dB at the leftmost point):

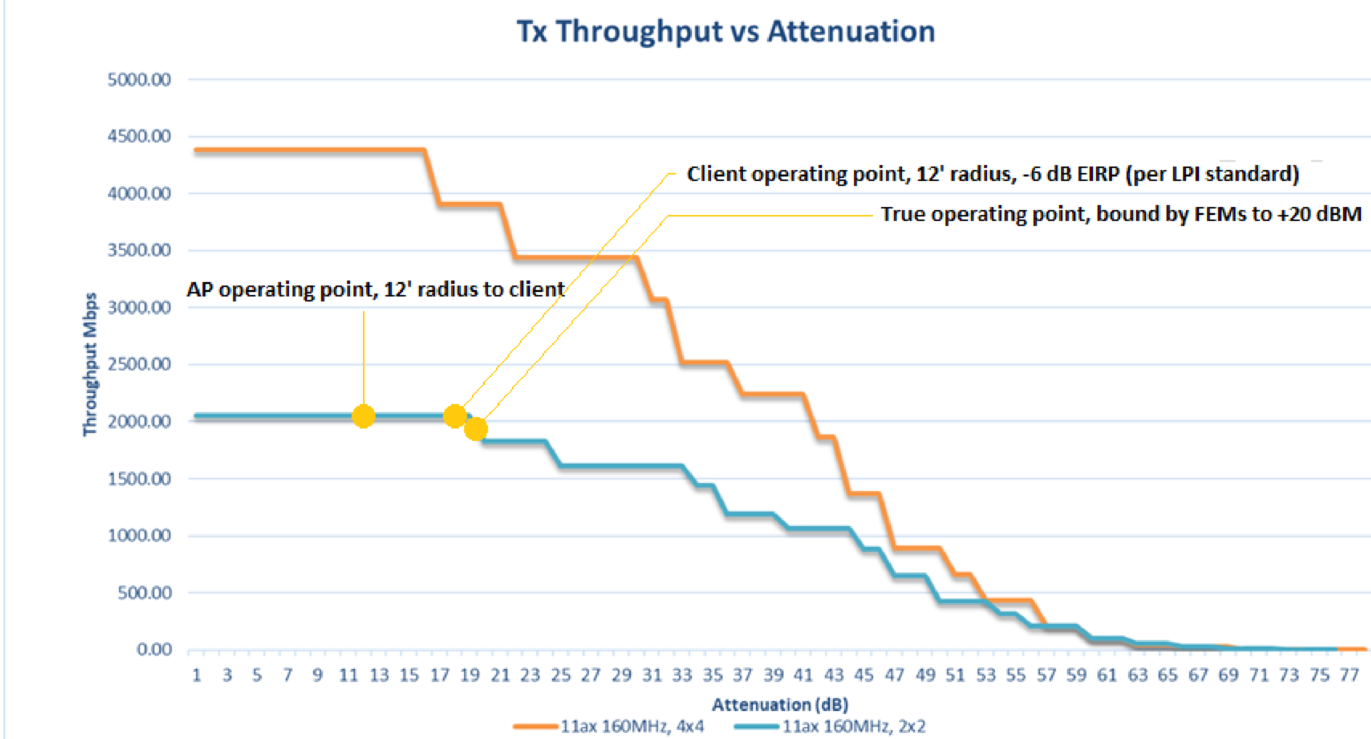

Figure 2 - Differential MCS Operating Points, Downlink and Uplink @ 12'

Note that the client – 6 or 7 dB down in power -- is operating at a mostly equivalent link budget to the AP at this close range, with a max compromise of one lower MCS step. This is not unusual behavior – and clients are almost universally the more restrictive on MCS setting than APs (certainly guaranteed when their transmit footprint is managed 6-7 dB below the AP).

2.3 Whole Home Performance

To get a feel for LPI coverage from an optimum AP location in the Wi-Fi house, the multi-band wireless router was placed in the central living room on the main floor. The mobile smartphone client was then moved room-to-room to measure data exchange rates and paint a coverage map for the house. Per housing floor, the TCP rates are shown below:

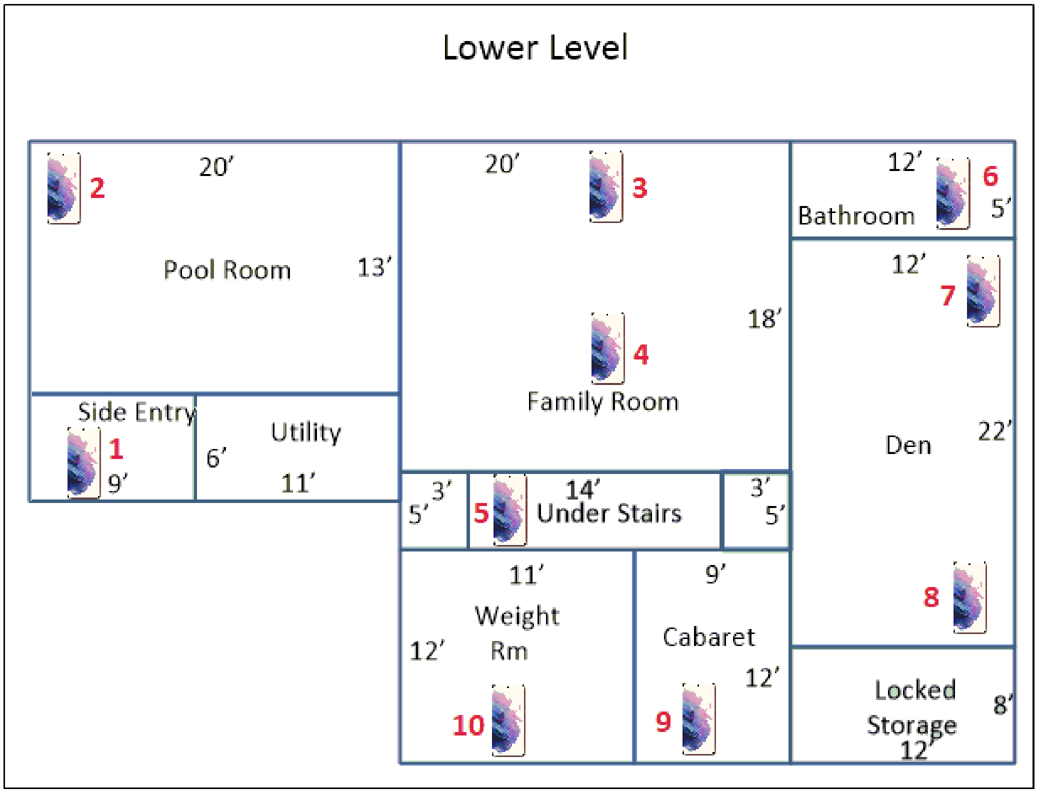

	TCP Mbps
1	480
2	650
3	1110
4	880
5	1280
6	540
7	660
8	540
9	830
10	970

Figure 3 - Wi-Fi Hous Lower Level 2x2 Client Performance

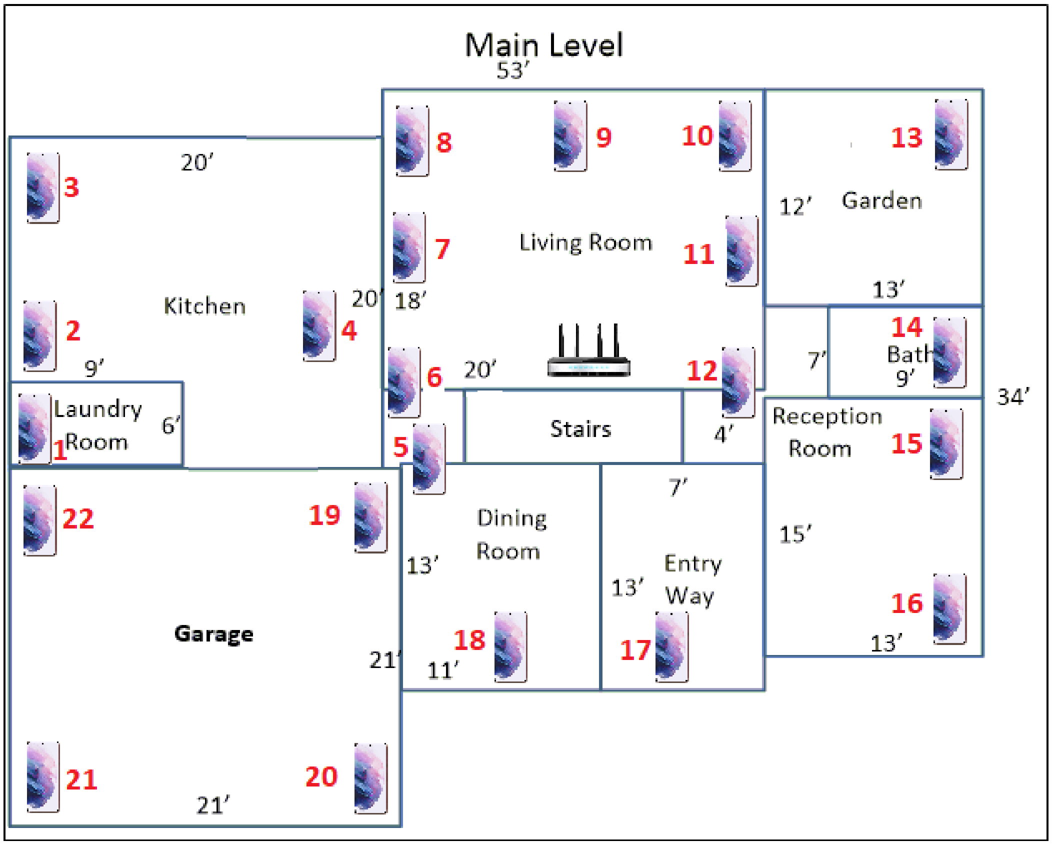

	TCP Mbps
1	1310
2	1140
3	1330
4	1470
5	1450
6	1480
7	1780
8	1610
9	1750
10	1440
11	1550
12	1600
13	1360
14	1450
15	680
16	780
17	1140
18	1320
19	1300
20	680
21	850
22	870

Figure 4 - Wi-Fi House Main Level 2x2 Client Performance

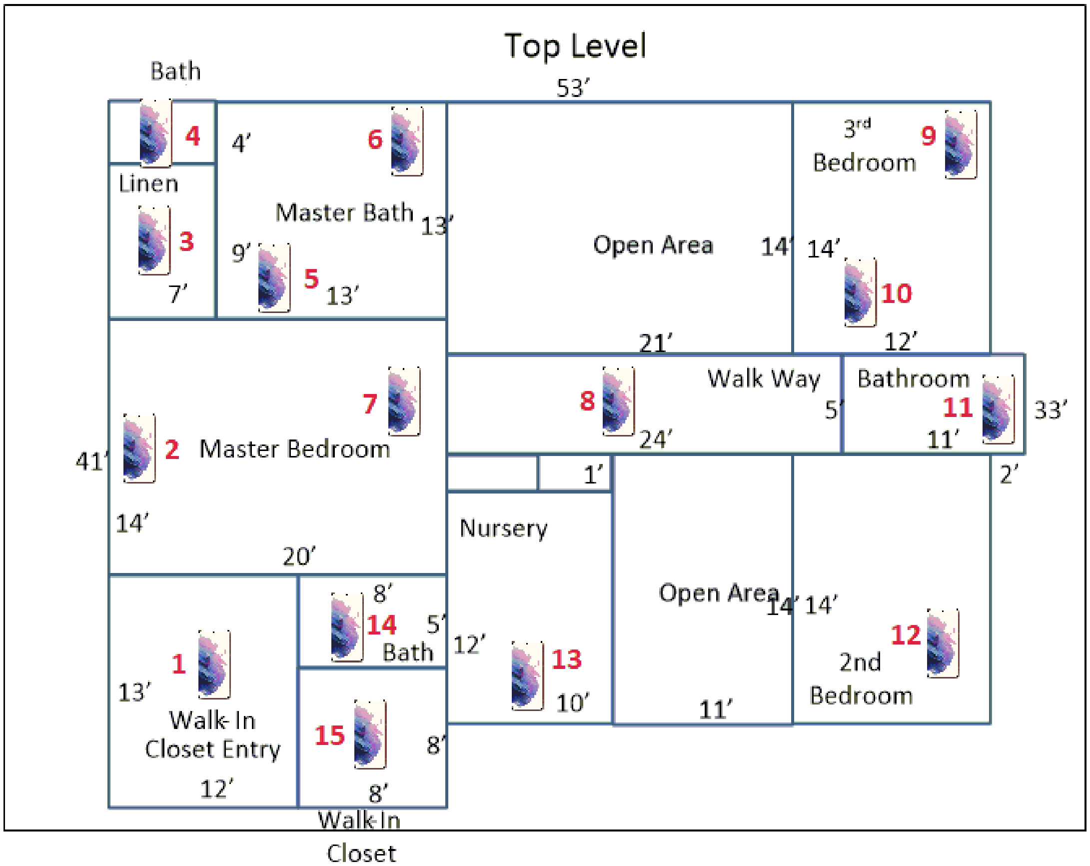

	TCP Mbps
1	200
2	540
3	690
4	960
5	670
6	820
7	650
8	1330
9	830
10	1140
11	960
12	650
13	830
14	530
15	300

Figure 5 - Wi-Fi House Upper Level 2x2 Client Performance

These collections of data yield the following “heat maps” of the LPI-backed off mobile smartphone around the whole of the Wi-Fi house:

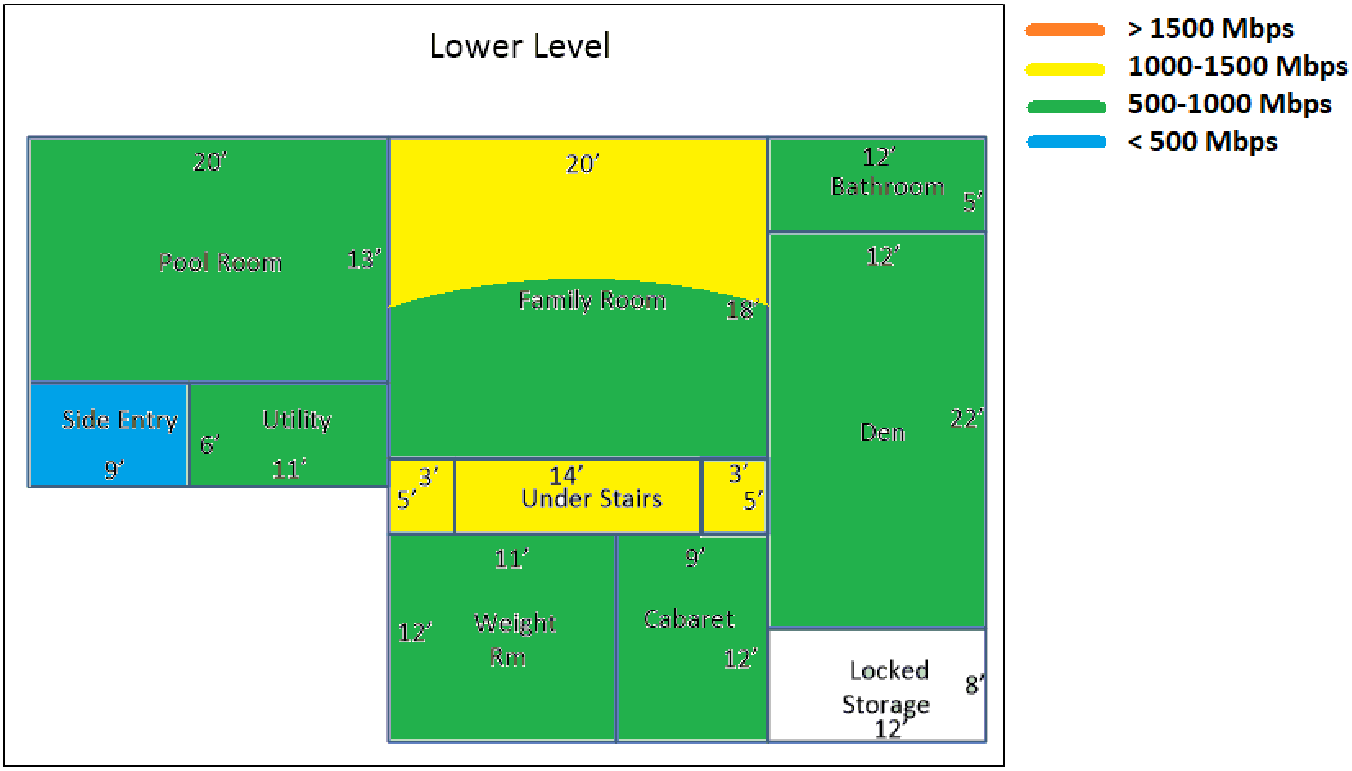

Figure 6 - Lower Level Heat Map

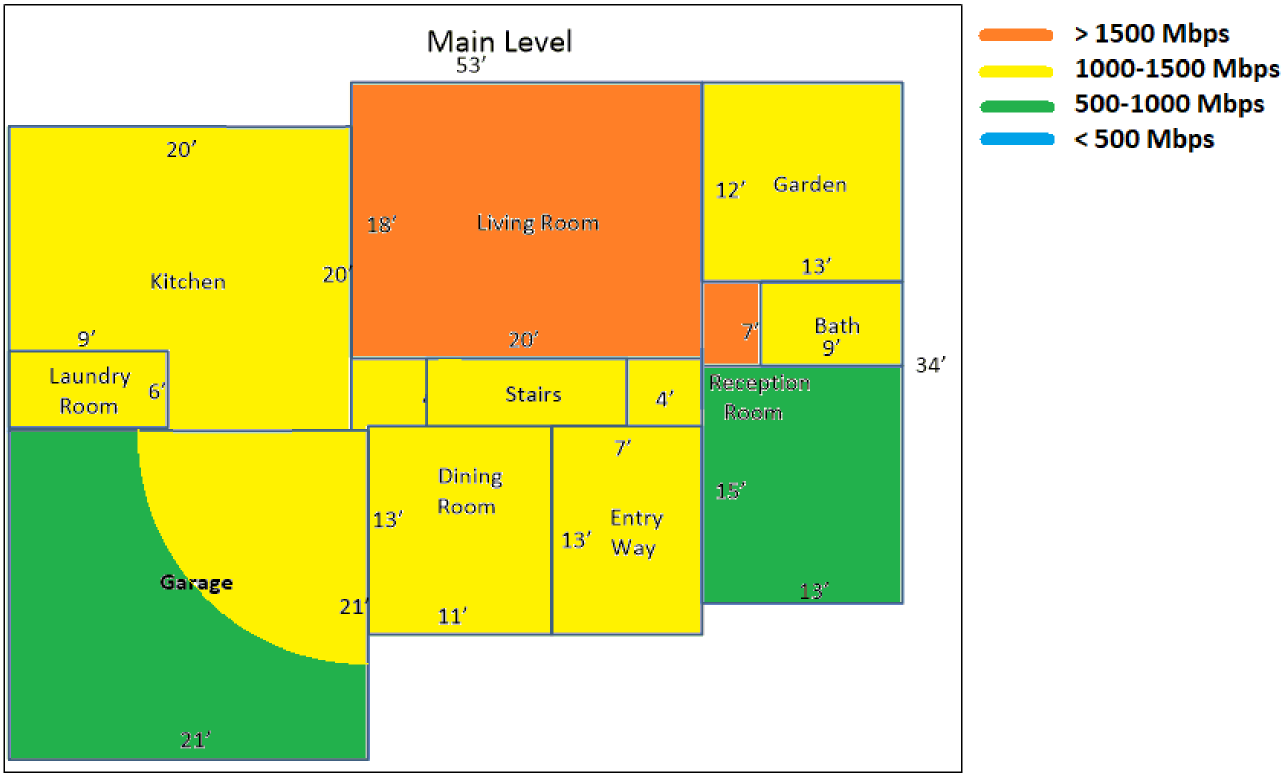

Figure 7 - Main Level Heat Map

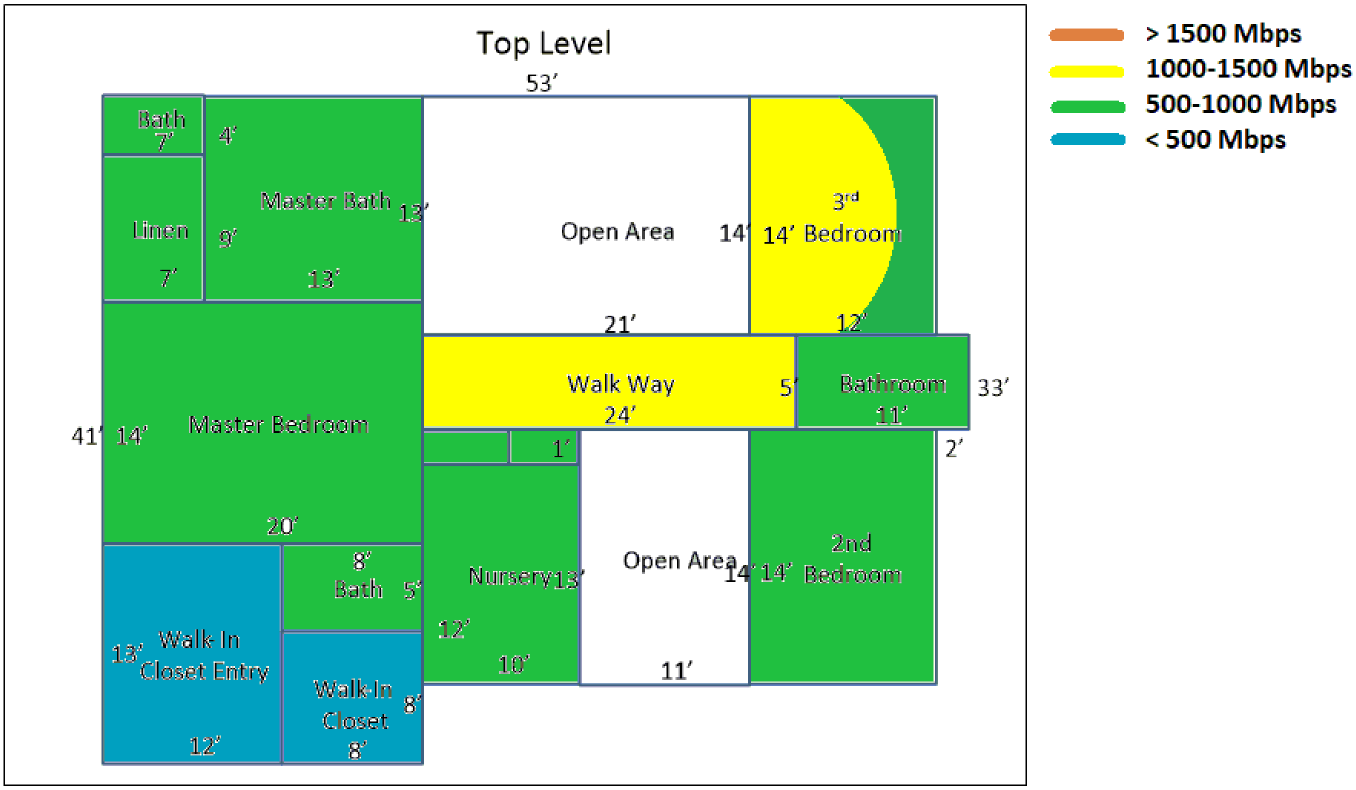

Figure 8 - Top Level Heat Map

Note that, on the lower level, there was a service null directly below the AP – likely a result of a compromised antenna gain pattern along the central vertical axis of that device. In any event, armed with these results, we can lay out expectations for coverage if we bump the AP power up to an AFC-approved 4W.

3. The Standard Power Prospectus

The fundamental tenet of higher EIRP is the notion that over identical path geometries, more transmit power means higher received power, better carrier-to-noise ratio (C/N) and better spectral efficiency (higher link MCS) which then sustains higher bitrate coverage over a greater service area.

3.1 Implications of the Power Gain

The bump from 27 dBm at 160 MHz BW in the 6 GHz band (LPI EIRP) up to 36 dBm is obviously significant. To put it in perspective, this implies a link linear throw increase of 2.8x to sustain at least the same MCS provided by LPI at a given service radius from client to selected AP or a service footprint up to almost 8x the coverage achieved at LPI for a selected reference bitrate – *provided the client can track the AP EIRP by no less than -6 dB*. (These are best-case numbers, essentially just FSPL without consideration for drywall and flooring – which, per transition, add lumped losses to the FSPL for a given radius). On the face of it, such an uptick suggests you can service a 10,000 square foot, multi-floor mansion from a single AP or gateway and achieve the same bitrate performance you would obtain throughout a 1500 square foot bungalow at LPI levels.

A couple of qualifying footnotes need to be appended at this juncture, however, to dampen and rationalize the more enthusiastic of these expectations.

3.2 Signal Propagation Considerations

The principal point of allowing standard power indoors (with an AFC system) is an acknowledgment that interior wall and flooring transitions will serve to attenuate and scatter 6 GHz Wi-Fi energy on its outbound journey. Further, depending upon exterior construction, building entry loss (BEL) anywhere from 6 to 30+ dB will constrain the radiation footprint on its inside-to-outside propagation. In prior testing at the Arris Wi-Fi house, multiple path tests indicated that, at mid-band 6 GHz, 4.5 dB drywall losses and 9 dB flooring losses could be inferred. (This will obviously vary somewhat with construction differences in other homes but seems a reasonable anchor in calculating available power at a given client location throughout the house, given the AP's own position and orientation.)

As a point of reference, it should be noted that these line-of-sight (LOS)-based calculations start to lose their predictive capability once you transition perhaps 3 walls/floors and get more than 40 feet from the AP – multipath effects not being accounted for (and one accumulates new signal energy vectors based on the dielectric thickness and size of every medium being traversed – as well as the sheer number of these transitions.) The best analysis involves a ray-tracing-based simulator (similar to the overlaid lighting effects summed in every computer-generated imagery (CGI) frame of a movie) but these can be expensive; lumped propagation channel models for in-home projection exist but require parameter tuning (related to the transition geometries associated with the main propagation path and the relative orientations of transmitter and receiver). All of the predictive work here is based on simple LOS analysis (to produce the most conservative link throw expectation); the bookmark is to just be aware of the sources of error in predicting more far-flung receive client behavior. Actual results at the range extrema would be better than our predictions.

Now, the setting of a goodput asymptote for a bitrate expectation versus path loss can overestimate the delivered signal footprint from AP to client if the spatial stream counts on both sides are equal. This is due to less-than-perfect antenna orthogonality and the effects of polarization mismatch (due to device orientation or path effects) between transmitter antenna farm and receiver antenna farm. In the cases where a 4x4 AP is lighting up various 2x2 clients, these diversity losses are washed out by what amounts to spatial oversampling available to the link by virtue of the spatial stream mismatch. However, this also means that the goodput expectations will scale to the least common denominator (2x2 in this case) – though the rate achieved will be closer to expectations in this scenario than for the case of a 4x4 device linking with another 4x4 device. (And all of this assumes fixed radiation patterns – i.e., no presumption of any type of antenna beam steering capability).

3.3 Qualifications for Standard Power Benefits

Using both the recorded LPI device performance at various in-home waypoints along with predictions for the same, and couple this with what amounts to a vernier or gain factor between expectation and measured performance at an open-look short link path, we can then begin to predict performance at standard power levels. Referring to our prior LPI measurements and projections at a small path loss in the Wi-Fi house yields the following observations: 1) Client spatial stream specification will define the rate multiplier (and hence, maximum asymptote) for delivered goodput bitrate performance; 2) Battery powered clients – with the constrained EIRP set by frontend module (FEM) thermal considerations (driving a more modest bias rail in pursuit of lower operating temperatures) – will not be able to track the increased output level of a standard power AP and hence these links will be compromised at large service radii as clients disconnect due to Wi-Fi framing loss; 3) AC-powered clients will benefit greatly from the power uptick, if they can attain a 30 dBm EIRP (24 dBm conducted with a maximum antenna gain of 6 dBi – which amounts to a much less demanding conducted power budget than that provided by the typical -1 to 2 dBi gains of small form factor client devices). Put another way, if portable client electronics can achieve a 20 dBm footprint with a 0 dBi antenna gain factor, then fixed AC clients can bridge 6 of the 10 dB power gap from tracking standard power client to portable levels (30 dBm down to 20 dBm) and establish the necessary standard power with only antenna changes and 4 dB worth of additional conducted power.

It is worthwhile to also note that these range considerations anticipate full leverage of the 160 MHz BW available at 6 GHz; the client connect limitations become dramatically more concerning if narrower BWs are employed on the links (leading to reduced client EIRPs – potentially down to 12 dBm if the channel operates at 20 MHz BW).

These observations produce an interesting set of architectural permutations of in-home Wi-Fi networks, as we shall see. Standard power goodput bitrate achievable for both 4x4 and 2x2 links is captured immediately below, along with operating points for the various uplink and downlink power budgets:

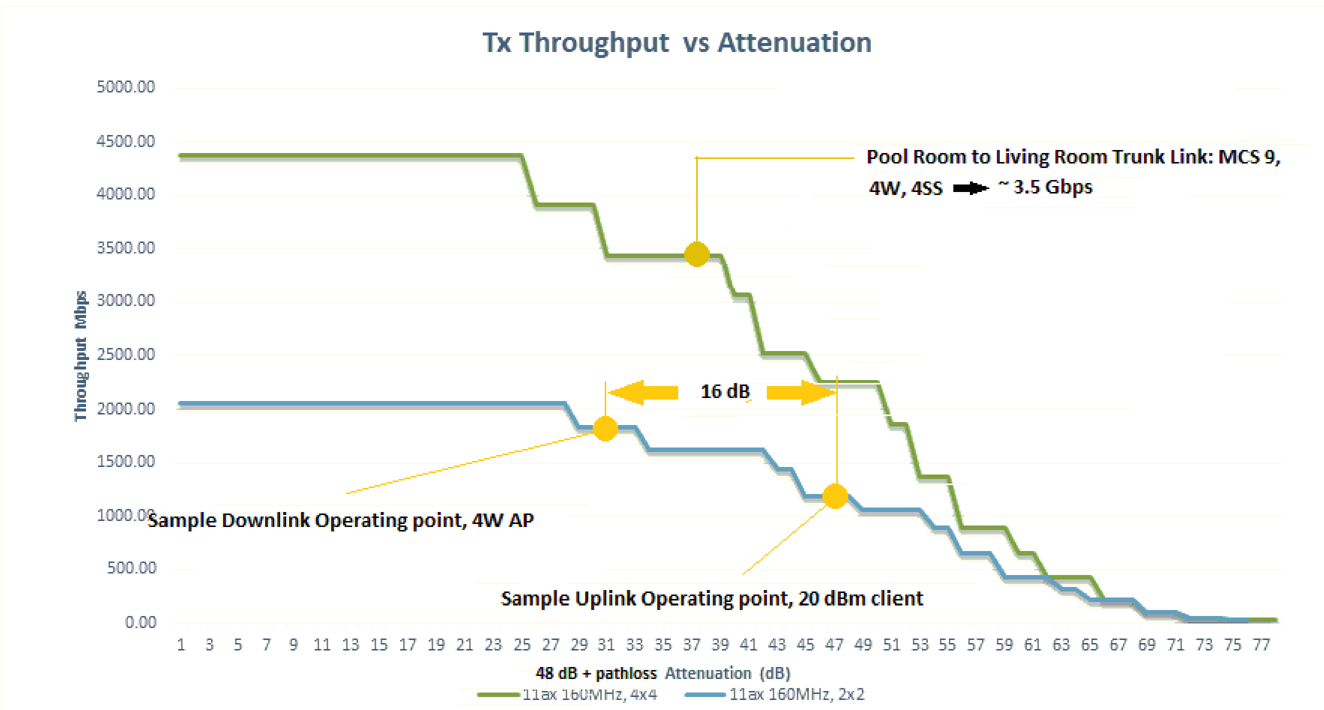

Figure 9 - Bitrate Waterfall vs Path Loss for 4W AP

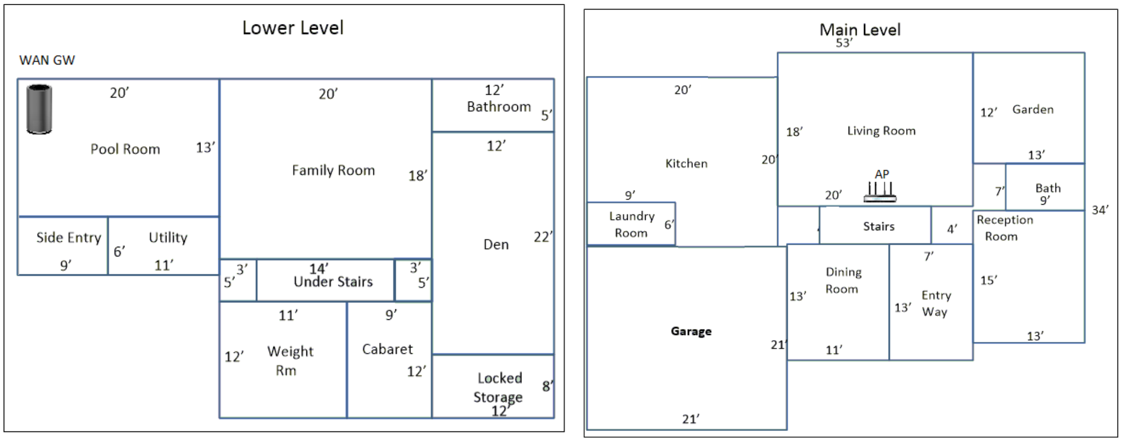

Figure 10 - Proposed Trunk Link Endpoints, Lower and Main Levels

As before, the ordinate axis assumes an antenna-launch loss of around 49 dB as a starting point and the values listed are essentially distance, in pathloss dB, from the antenna(e) to the client. Fundamentally, the move to higher power represents a push (reach expansion) on the pathloss axis by 9 dB:

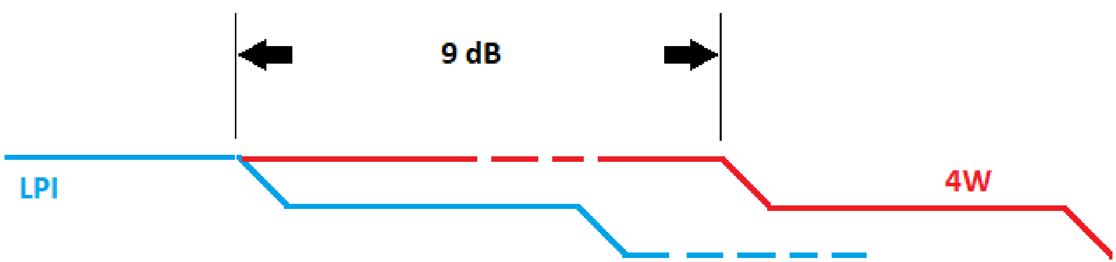

Figure 11 - Effect of Increasing EIRP on Delivered Bitrate

The constrained radiation footprint of portable, battery-powered client devices puts a damper on the more inflated of the coverage aspirations one would otherwise assign to a 4W EIRP, servicing AP, then. But the more robust (4x4) trunk performance and ability to reach more 1W-capable (largely fixed location) clients remains. The question arises: at what size in home floorplan would one want to invest in 4W APs, and what strategies on LPI and 4W mix might make sense? These considerations for in-home wireless service mounts are addressed in detail in section 7. But first, the machinery of AFC, its motivation and operational implications need to be understood (if only as background detail).

4. Potential for Interference

The operational "hall pass" granted Wi-Fi exploit of the 6 GHz spectrum at up to 4W EIRP levels comes (rightfully) with accommodation for exploits of that spectrum by existing communications infrastructure. It is in this accounting and policing of the asset by the FCC that we can arrive at fair use of 6 GHz, so it is worthwhile examining the known fixed wireless access (FWA) actors in the space.

4.1 The Contenders for Common 6 GHz Spectrum

While a reasonable directive for oversubscribed spectrum amounts to "listen before talk", the fact remains that simultaneous access to the same piece of the 6 GHz band for different unlicensed communication systems would be problematic absent some appreciation (and interdiction) of the opportunities for one system's transmissions to compromise the receptions required of an alternate system. The spectral region of concern here involves the 850 MHz worth of U-NII-5 and U-NII-7, where the FCC is now granting Wi-Fi the opportunity to extend its reach by leverage of AP power up to a 4W footprint. The contending system cases can be shown as follows:

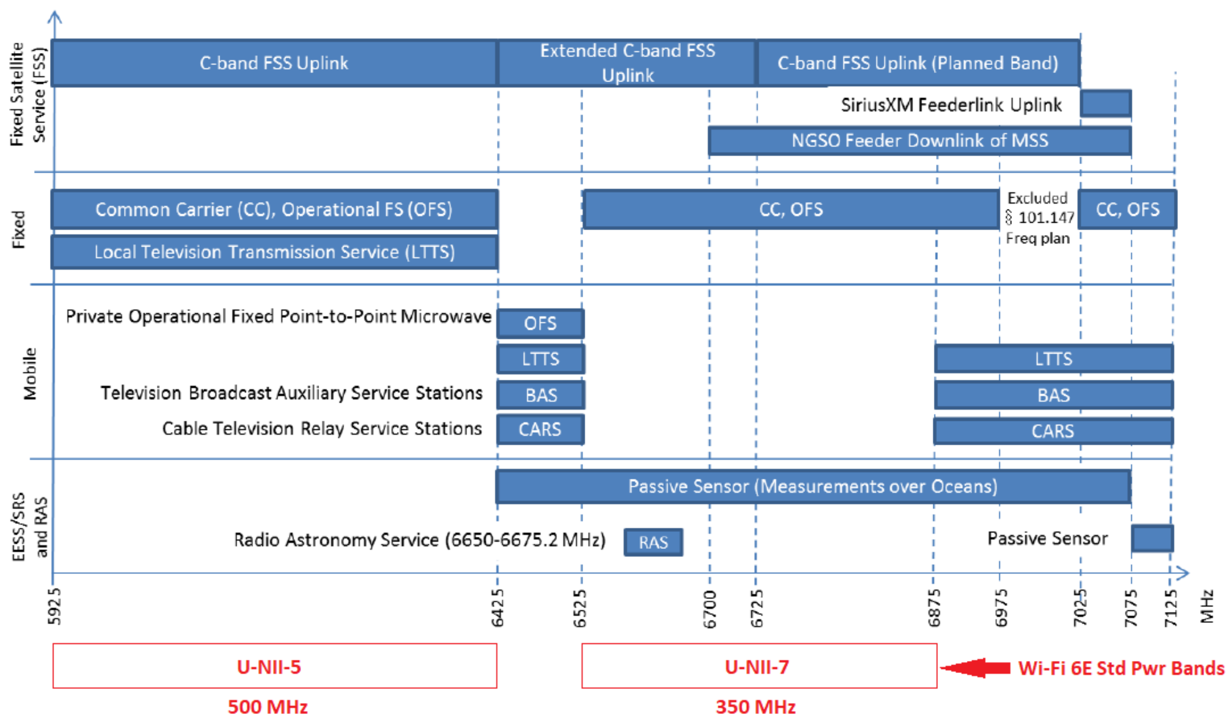

Figure 12 - Legacy FWA 6 GHz Spectrum Use

Note that the services described encompass both fixed service (FS) and C-band fixed satellite service (FSS) mounts, along with some mobile point-to-point (P2P) links as might be employed by news services with "on the scene" coverage. Given the predominant population of FS clients (which are also increasing in number, where C-band satellite growth shows a -2%/year reduction) and the fact that the much smaller number of mobile services (MS) have the ability to re-orient antennae to mitigate interference at the deployment site, we will focus on the impact to FS infrastructure. However, low-horizon satellite positioning (in the norther latitudes) can potentially expose the satellite receivers to rogue, unlicensed ground link propagation paths at low elevation offsets (presenting much less than orthogonal look angles) and so an effort should be made in establishing the relative heft of 6 GHz indoor energy impinging on the geosynchronous equatorial orbits (GEOs) used by C-band satellites (whose uplink band overlays U-NII-5 and -7 as shown above).

4.2 Dismissing Satellite Interference Concerns

In the report *Frequency Sharing for Radio Local Area Networks in the 6 GHz band, revision 3,* prepared by RKF Engineering Solutions for the 6USC, there is an exhaustive simulation assigning 6 GHz devices (as either small cell or Wi-Fi devices) on a per-person basis, with historical representations on device use cases which serve to estimate the concurrent, accumulated radiation from all indoor standard power devices across the continental United States (CONUS) and estimate the interference-to-noise ratio (I/N) as presented to the GEO-parked C-band satellites in use. Fundamentally, duty cycles of use on high power devices, based upon inferred link capacities and data consumption per hour were assigned. The upshot of the protracted chain of assumptions was that ~ 1 billion devices could be involved across CONUS, with the instantaneous overlay of around 400,000 standard power ON cycles. (The report spells out all assumptions used in the simulation and the reader is invited to download the paper – which is available on the internet in the public domain.) For comparative purposes, the report also estimated the impact to satellite receivers of the legacy 6 GHz infrastructure (which obviously is already operating concurrently with C-band satellite without measurable negative findings). A key figure of the study shows the following impact to receivers on the GEO arc from both legacy systems and the expected high indoor use of 6 GHz standard power Wi-Fi:

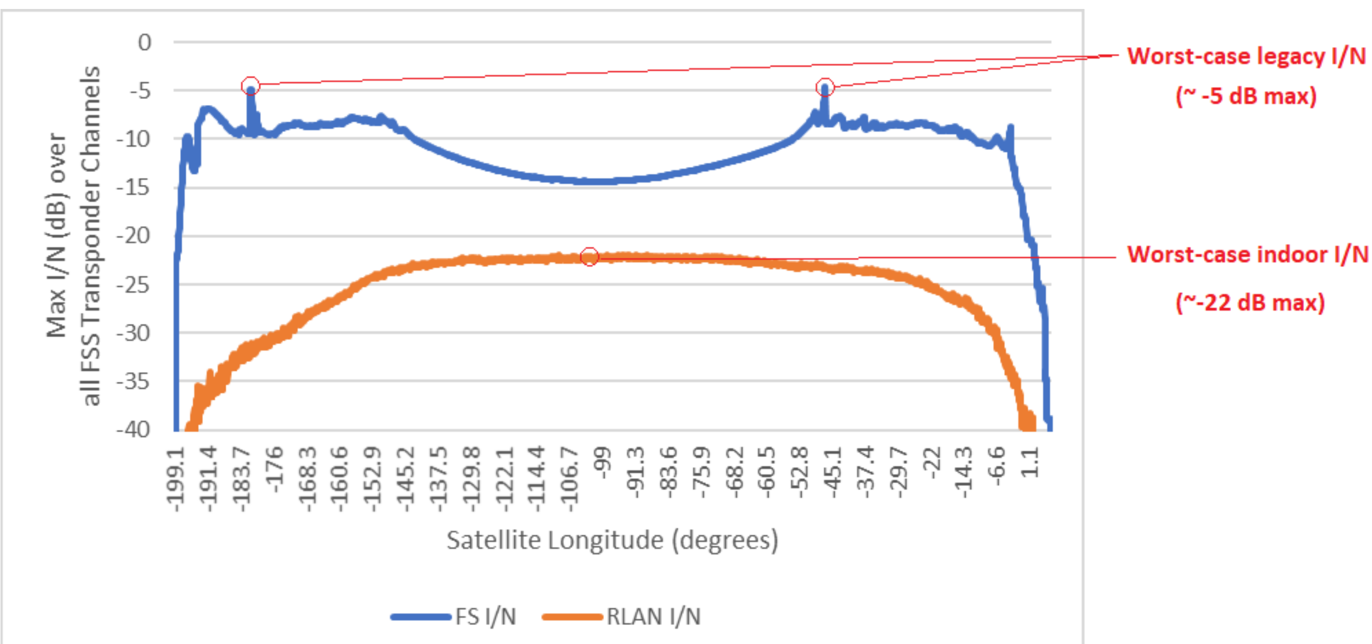

Figure 13 - Relative Obscurity of Potential 6 GHz Indoor Wi-Fi Interference to Satellites on the GEO Arc

The estimated 17 dB lower I/N due to indoor standard power Wi-Fi – versus what the satellites experience already from legacy 6 GHz outdoor links – is a telling graphic (and allows us to concentrate on interference implications to land-based FS links only). This consideration summons AFC exclusion zone calculation around these legacy endpoints and as such, will be addressed in the following section on AFC.

5. Band Management: Automated Frequency Coordination

5.1 The Premise

The intention of AFC to create an FCC-supervised gate into unlicensed indoor exploit of 6 GHz begs elaboration. In order to instantiate this higher performance gate, a specification has been developed which merges several databases, an exclusion zone calculator, the presumption of default LPI behavior by APs seeking higher EIRP, a restriction on indoor-only AP use and a messaging exchange function (between a cloud supervisory portal and those APs). The block diagram manifests as follows:

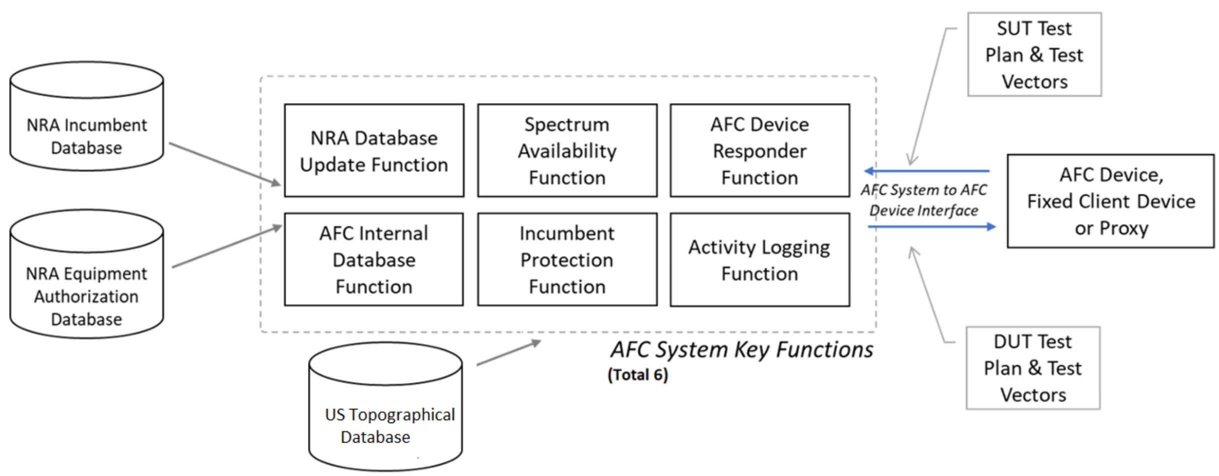

Figure 14 - AFC Block Diagram

As the diagram indicates, there have been six System Key Functions identified for the specification-in-progress, along with three reference databases. The intention of this organization is to support a bilateral data exchange with AP's soliciting the AFC for permission to operate at standard power; by and large, this exchange consists of a query from the AP which provides its latitude, longitude and operating height information to the AFC, which then references its databases for identification of the location of potential co-channel interference (CCI) targets and delivers a mask of available band and power back to the AP. The abbreviated schematic also points the way to certification – system under test (SUT) and device under test (DUT) test vectors at the point of data exchange (interface to the right in the above figure).

Key to all this is the presumption that, unless the appropriate provisioning message (parameterized permission) is received by the AP, its default operating EIRP envelope will be constrained to the LPI-specified 5 dBm/MHz PSD. Furthermore, the permission assigns a lease timeout period after which even approved standard power devices are required to revert to an LPI emissions envelope. (This can obviously be forestalled by periodic queries to the cloud AFC, which would then serve as renewals of the standard power lease). The buried requirement here is that the AP must be able to access the AFC portal at least once per day (the currently mandated expiration timeout) – and must repeatedly guarantee its lease renewal in order to maintain permission to operate at a standard power envelope.

As of this writing, many implementation details for AFC remain unsettled (key among these being the hosting question for the application – which then begs additional work on whether the service is industry-sponsored or a private affair -- the latter, then, invoking appropriate subscription or licensing). As might be expected, multiple stakeholders in the venture are currently deeply involved in finalizing details, since the goal is to have an operational AFC some time in 2022.

5.2 Operational behavior

The following outlines the specific functional parsing associated with the cloud-based AFC (pertinent to block diagram above):

Architecture/Function Parsing

- NRA (Nat'l Regulatory Authority) Database Update Function
  - DB of incumbent links w/locations, descriptors and credentials (maintained)
- AFC Device Responder Function
  - Duplex cloud link (URL based) which provides HTTPS/JSON portal for AFC device comms
- Spectrum Availability Function
  - Generates payload for response messages to devices (incl. error msgs)
  - Invokes Incumbent Protection Function and Logging Function
- Incumbent Protection Function
  - Math engine to do interference calculations (both CCI and adjacents) and recommend permissible channels (and operating power levels)
- Logging Function
  - Creates/maintains "non-repudiable ledger" of AFC transactions
- AFC Internal DB Function
  - Largely parametric details on incumbent installations (as antenna pattern specs and related)

Figure 15 - System Key Function Responsibilities

Message exchange details provided below:

- Northbound (device to cloud)
  - Available Spectrum Inquiry Request
    - Unique ID
    - Device Descriptors
    - Location Detail
    - Inquired Freq Range (MHz) and/or
    - Inquired Channel Numbers
    - Minimum Desired Power (dBm or dBm/MHz)
    - Vendor Extensions

- Southbound (cloud to device)
  - Available Spectrum Inquiry Response
    - Unique ID (per upstream request)
    - Allowable PSD by Freq Range (dBm/MHz) and/or
    - EIRP by List of Channels
    - Expiration time for provided ops (GMT)
    - Response Codes (P/F and error codes)*
    - Vendor Extensions

*Pass/Fail, with codes 100-199 being reserved for errors related to message formation, authentication, etc and 300-399 for tech editing concerns (like requesting inappropriate/wrong channels)

Figure 16 - Messaging Details Between AP and Cloud AFC

Additional messaging clarification is provided by the following:

- Device Descriptor is a 3-tuple: serial #, FCC ID# and (for US) a text string “47_CFR_PART_15_SUBPART_E” (would be different for other countries)
- Location is longitude, latitude and height (as degrees relative to the central meridian, degrees relative to the equator and meters above local terrain). The location footprint of the AP(s) in question may be expressed as an ellipse or 1 of 2 versions of a polygon area. Uncertainty self-certified (but reported) and an enumerated field describes whether the unit is indoor or out.
- Inquired Freq Range is as “a-b” where a, b are in MHz
- Inquired Channels is an explicit list of requested channel numbers

Figure 17 - Messaging Supplemental Definitions

Note that the exchanges listed may be executed by a proxy device (mobile, for example) on behalf of the AP. This would be accomplished with an NFC link between AP and proxy which would facilitate the cloud link connection and grant a pathway for user options and supplemental data entry on device location. The proxy shim would appear as follows:

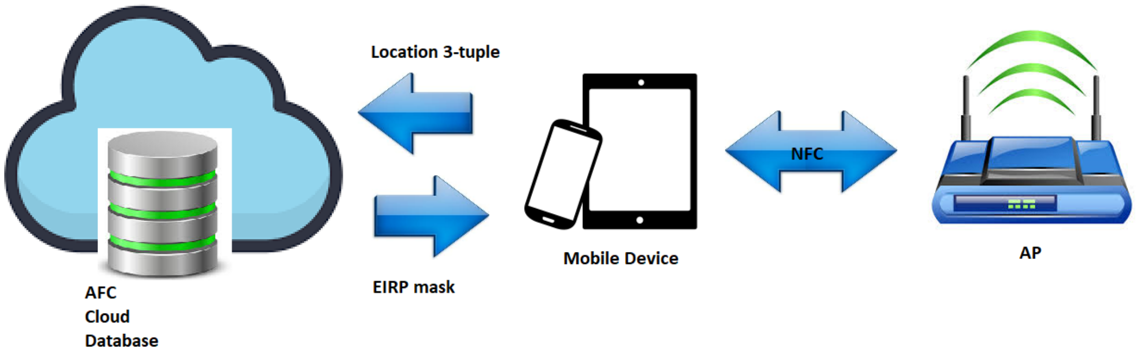

Figure 18 - AP Proxy Arrangement

The motivation for this added proxy lies in leveraging various methodologies to extend GPS position indoors, away from building apertures such as windows and to the immediate proximity of the AP (which, in a buried interior placement, may otherwise be blanked for GPS coverage). It also provides an application hosting environment where additional location data (height in floors or feet) may be inserted via the user (to improve on three-dimensional (3D) AP location uncertainty).

5.3 Nuances and Potential Optimizations

The operational dynamics of the AFC system suggest that, in the interest of speedier resolution of atomic requests from APs for access to standard power operation and a reduction in the query traffic impinging on the AFC from a nationwide population of devices, some amount of "query triage" might be considered. The driving consideration in this is that for a vast number of the requesting APs, lack of proximity to fixed wireless 6 GHz links guarantees deep marginalization of CCI possibilities which might otherwise negatively impact operation of the legacy fixed wireless links. In fact, CCI requires a multilateral conspiracy of overlaid/adjacent channels in the band, aligned FWA antenna aperture (as both azimuth and elevation above ground) and minimal geographical offset of the two contending endpoints in order for concern over potential CCI to be realized. These observations present the opportunity to streamline determination of exclusion zone limitations (reduce calculus overheads) and perhaps extend standard power AP leases past the 24-hour limit presently in the specification (in cases where the potential for interference is effectively and persistently nonexistent). Both efficiency gambits invoke some subtle tradeoff considerations, however.

5.3.1 Radius of Interference

First, as regards the reduction of calculation overhead, is the observation made above that the potential for CCI involves competitive channel access, unfortunate alignment of the FWA antenna with the source of possible interference and geographical proximity of the two potentially interfering endpoints (legacy FWA and Wi-Fi AP). Of these three parameters, inter-endpoint distance is easily the most significant contributor to CCI. It is only when this radius falls below some minimum that additional calculation need be performed to ascertain the possibility of CCI at a level which threatens operation of the FS or FSS link. For example, free-air LOS path loss over the radius involved (assuming optimal coupling of antenna apertures) and a conservative assignment of only a 10 dB penalty for dwelling egress (likely values will fall in a range of 10-30 dB at 6 GHz, depending upon exterior wall construction details and proximity of a standard – non-e-glass -- window) generate the first cut at the potential for CCI based solely upon proximity of the FS receiver to the Wi-Fi AP.

Some example numbers would be helpful: at 6.5 GHz (mid-band 6 GHz) free-air path loss amounts to 113 dB at one mile. Coupling this with the maximum indoor AP power of 36 dBm and a building entry loss (BEL) of 10 dB suggests the propagated Wi-Fi signal amounts to no more than -83 dBm at a one-mile radius in any direction. The goal, however, is to get to I/N to -6 dB at the nearest FWA antenna (i.e., no closer than 6 dB below the noise floor at a particular BW). To be especially conservative on this gambit, we will neglect the interference moderating effects of clutter, scatter, terrain blanking (available in several standard outdoor propagation models) and narrow antenna aperture (assuming, in this case, no azimuth selectivity in the FWA antenna systems). For a 160 MHz overlaid band (between FS and AP) and a presumed 4 dB NF on the FS receiver, we get a maximum allowable interference footprint of -174 + 10*log(BW) + 4dB receiver NF – 6 dB minimum desired I/N ratio – 33 dB (FS boresight antenna gain for a small dish), or -127 dBm. Translating this to path loss of the AP signal (and accounting for benefitting only 10 dB on the BEL), the distance would have to amount to +36 – (-127) – 10, or 153 dB. At mid-band, this free-space loss comes at 102 miles displacement! Clearly, a single distance-to-FS receiver metric (under these stacked, overly conservative presumptions) does not represent much exclusion zone calculation triage benefit.

Perhaps a more realistic approach would be to find the distance at which a 4W AP presents no worse an interference candidate than an LPI device (which is permitted an indoor EIRP of 27 dBm @ 160 MHz BW *anywhere* in the US) – and then append a keep out area which, at its periphery, represents no greater a threat to FS operations than the energy footprint of an indoor LPI device exiting its enclosure (to a total throw of 300 feet or so) through an especially stout BEL. To set that energy threshold, one can conservatively set the BEL for the building housing the LPI AP to 30 dB (assumes stone exterior construction and e-glass windows), add 300 feet of arbitrary FSPL (as nearest proximity to an FS endpoint) and then use this abated field strength as the target value for a 4W AP to hit (under a less beneficial BEL of 10 dB, just to bake in margin). In round numbers, the target signal strength ends up being 27dBm – 88 dB (FSPL for 300 feet, mid-band) – 30 dB (BEL), or -91 dBm. This implies that one requires the FS spacing to the standard power AP to account for 117 dB worth of path loss (36 dBm – 10dB BEL – (-91 dBm target)). This arrives at a throw of 8500 feet (1.6 mi) at mid-band 6 GHz. This makes for a more manageable first cut.

So now the first triage step is to determine if the AP is further away from the nearest FS receiver by at least 1.6 miles; if this is the case, then the entirety of the U-NII-5 and U-NII-7 bands could be released for use by the petitioning AP – and the concern for interference potential relegated to being no worse (and this, actually by a fair margin) than that supplied by an equivalent indoor LPI device located in a stone home 100 yards from the FS receiver.

5.3.2 Antennae Alignment

Now the notion of antenna selectivity can be introduced. Note that with the 1.6-mile approximation above, all signal loss was put down to FSPL and BEL. With one additional step, APs which locate themselves less than 1.6 miles from a potential FWA site may be cleared to operate at standard power by evaluating whether the combination of FSPL, BEL and antenna pattern misalignment could defuse the possibility of interference. Note that FS endpoints have antenna radiation parameters logged into the Universal Licensing System (ULS, also referred to as national regulatory authority (NRA) in separate documentation) database referenced by the AFC. For the prevalent class A and B antennae, the pattern is fairly narrow around the boresight direction of the element (for both azimuth and elevation). (In situations where sectorized coverages are combined – as would be the case for point-to-multipoint (P2MP) systems -- it is straightforward to select the most aligned element(s) of the array). An example of the most common apertures deployed in the field is captured below:

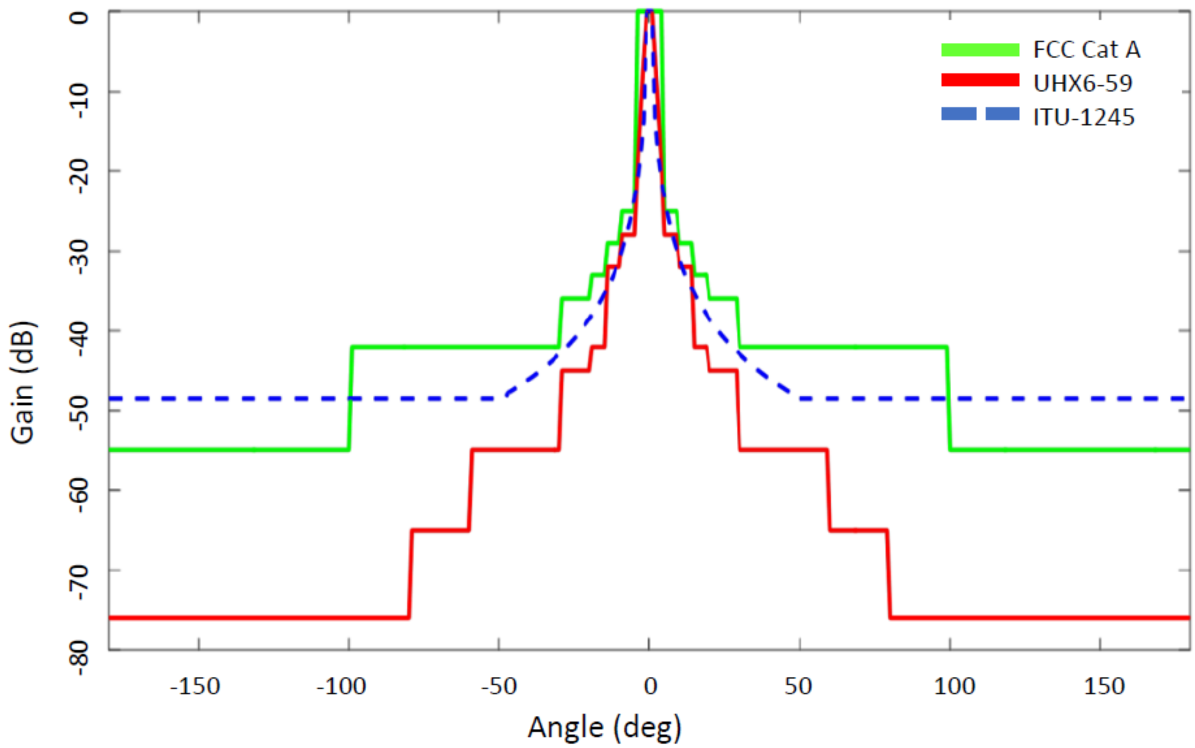

Figure 19 - FS Antenna Apertures (Azimuth Look Angle)

Zoomed in, the aperture looks like this:

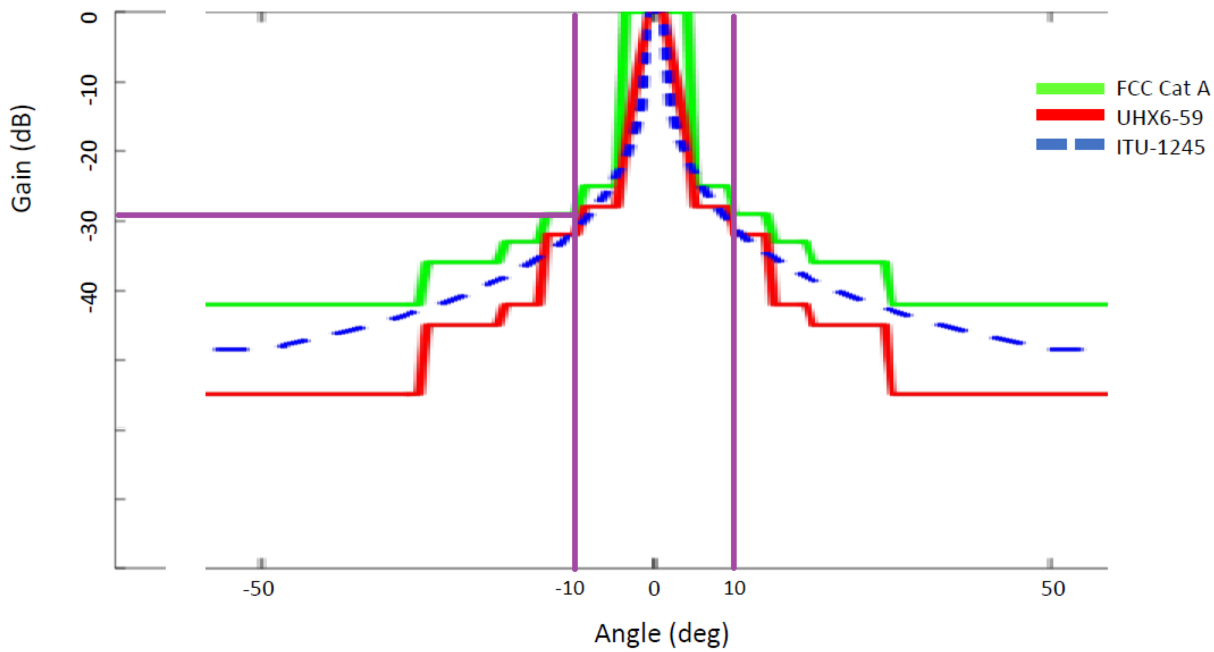

Figure 20 - Antenna aperture zoomed

Note that such performance is representative of 83% of the FS antenna population (per Commsearch research in 2011 – such not being recently re-evaluated, however). Even a cursory examination of the figure reveals that, once off-boresight by as little as +/- 10 degrees, any of the commonly used FWA antenna masks shown exhibit an azimuth selectivity which will produce > 30 dB loss relative to the aligned gain. Normally, one would consider boresight gain as the reference against which, would be applied the selectivity offset loss; however, recall that the distance calculation above applies to any angular displacement of AP-to-FS receiver, including exact boresight. Referencing our prior interference radius calculation, if we can count on an *additional* interfering signal loss due to the FS receive antenna's look angle to the candidate AP, our FSPL budget reduces to 117 – 30, or 87 dB. In mid-band, this would allow Wi-Fi AP's full rein on standard power as long as the AP distance from the FS endpoint is greater than 270 feet and the look angle exceeds 10 degrees on either side of target system antenna bore sighting. But (though in a significant minority) there are other antennae in use by legacy 6 GHz FWA infrastructure which feature a broader aperture and we will consider them, despite a low representation in the total number of target antenna systems, as the least common denominator default in calculating a conservative exposure profile for nearby FS receivers.

A representative sample of a popular (smaller) 0.9m diameter antennae, the Andrew VHLP3-6W, exhibits the following selectivity for horizontal and vertical, single polarizations:

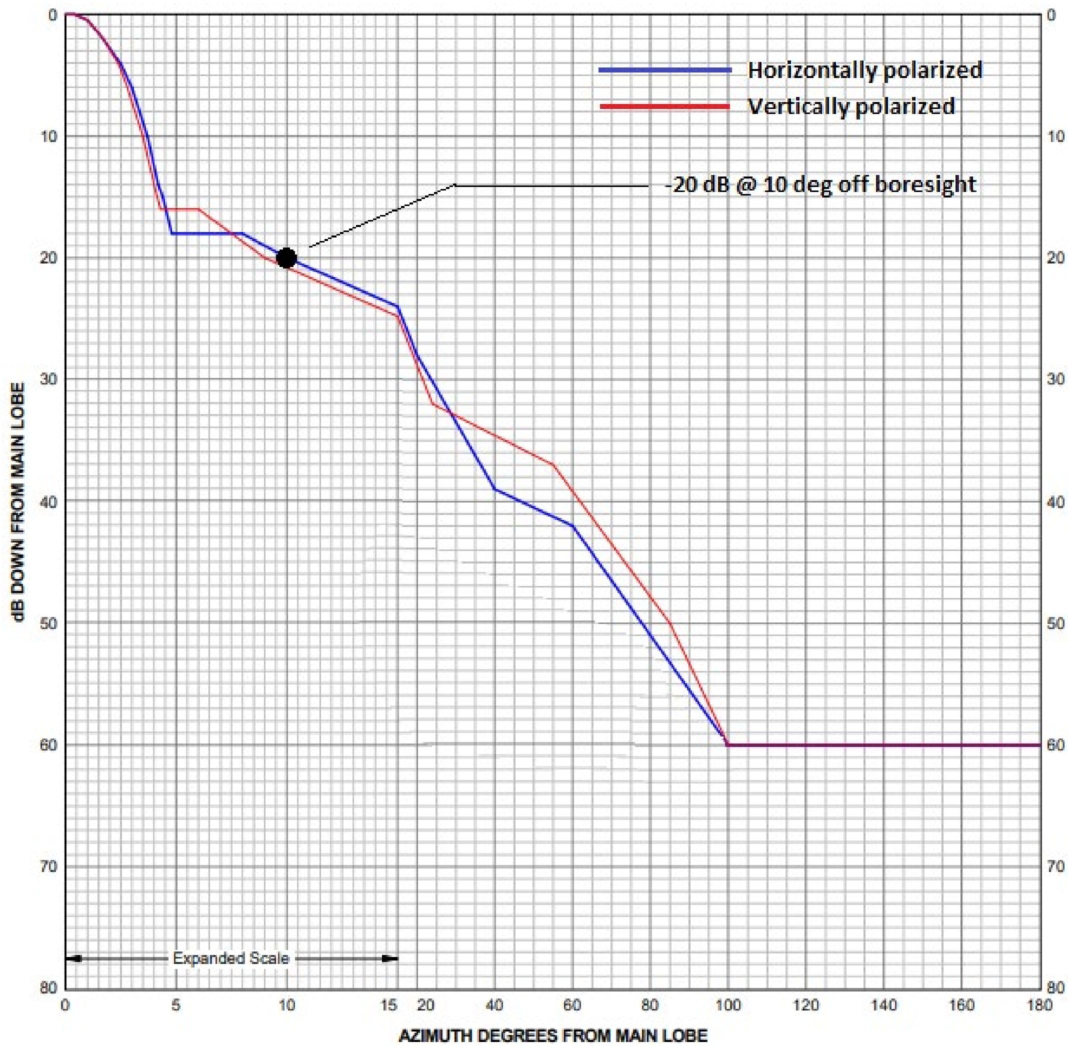

Figure 21 - Representative Small Diameter Antenna Single-Sided Selectivity @ 6 GHz

If we peg look angles > 10 degrees (between AP and FWA antenna) as a reference point, this smaller antenna exhibits worse selectivity than larger units (only ~ -20 dB or better, relative to the main lobe – versus the -30 dB exhibited by the majority of the FS antenna population). Using this as our generic interference possibility, our margin to interference collapses to 117 – 20, or 97 dB of required FSPL. The mid-band proximity associated with this number amounts to 850 feet (0.16 miles). The second "shortcut", then, for releasing the full extent of U-NII-5 and U-NII-7 spectrum for standard power use of APs has the dual qualifications that the distance to FS antenna must be > 0.16 miles *and* be off-boresight by at least 10 degrees.

Outside of these calculation shortcut cases, the AFC system would likely have to perform a full exclusion zone calculation (using distance, look angles, FS receiver link details and related) to evaluate whether Wi-Fi interference is a consideration.

5.3.3 The Argument for Variable Lease Duration

As regards the notion of extended leases, the FCC has decided that 24 hours (with an additional 24 hours grace period in the event of AFC system connection failure) will be the norm. This is based on the mere possibility of the ULS database being amended on 24-hour boundaries. Originally, the 6USC consortium had proposed a 30-day lease, projecting that the potential for interference nationwide would not significantly alter within a one-month timespan. It seems very likely that, in rural and suburban areas, the pace of 6 GHz infrastructure build-out will see a rate of database change very much slower than every 24 hours. It may be worthwhile for the FCC to re-examine standard power lease epochs – if for no other reason than this would greatly down-scale the communications traffic from APs to the AFC system (though, in fairness, this economy of communication does not appear to be an impediment for the system to operate, given the vast number of network paths in place to gather the AP solicitations and the ability to map AFC from the single cloud portal to multiple network edge locations). There certainly seems to be a case for applying some variation in lease lengths, depending upon location parametrics associated with the petitioning APs – especially if a case can be made to the AFC system that the device location is fixed (i.e., not a portable device).

Referring to our section on AFC calculation triage, a great many of the soliciting APs (perhaps as many as 90%), should require no full calculation of interference potential – which suggests longer leases may serve adequately.

5.3.4 Apportioning the Shared Spectrum

The final option we consider here is that, if the probability of the two wireless systems interfering with each other is not easily dismissible based upon proximity and alignment, we can examine the leverage of alternate in-band spectrum (non-overlaid assignment of spectral resources) to resolve the issue (and so grant some amount of leeway to the requesting AP(s), based upon available contingency spectrum). By appending guard bandwidth on either side of the incumbent system's operational link bandwidth, an exclusion portion of the spectrum for the soliciting Wi-Fi AP may be determined. (This would have to be done for every exposed FS node within a 1.6-mile radius -- depending upon boresight offset – each potentially with a differing link spectral profile.) Note that legacy FS links typically feature system BW << 100 MHz and there is 850 MHz of standard power 6 GHz spectrum across the two leverageable U-NII bands which can provide the frequency-division multiplexing (FDM) options desired – so the opportunity for mutual avoidance is quite high in all but the most crowded urban environments. Further, the Wi-Fi APs are already required to meet beyond-band-edge transmit power rolloffs which conform to the following out-of-band (OOB) energy mask:

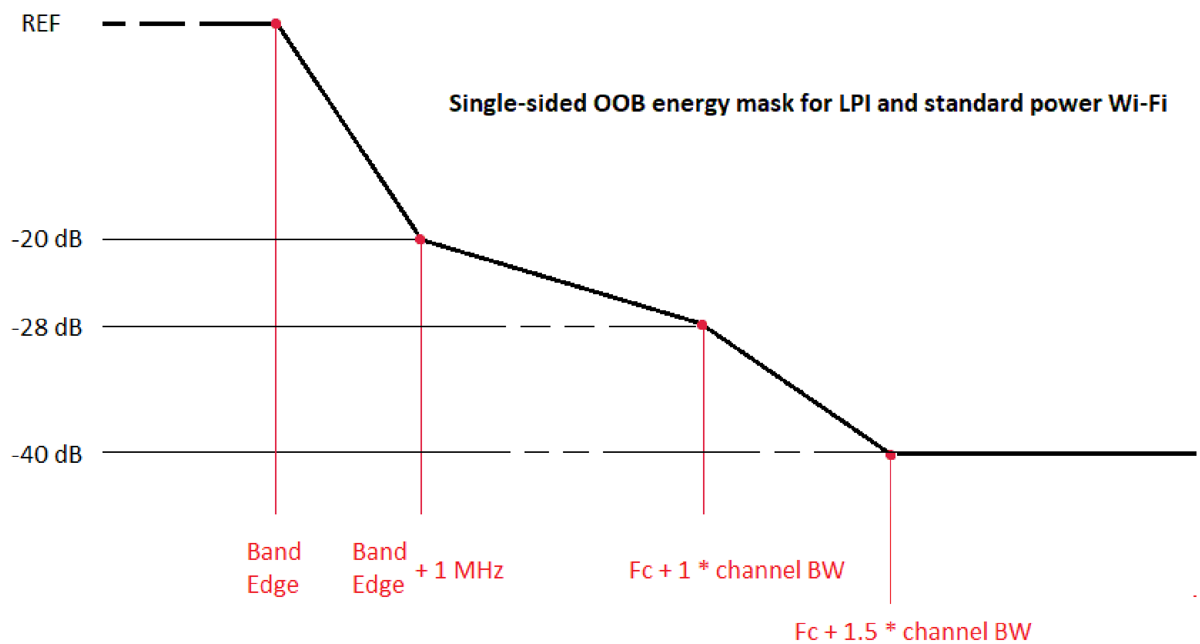

Figure 22 - Required OOB Energy Suppression by Wi-Fi APs

With 20 dB of specified OOB suppression at +/- one MHz from either band-edge of the AP's channel, and presuming 160 MHz of AP BW to set guard bands, one is tempted to represent the AP's spectral footprint as 162 MHz and have that augmented edge-of-band be no closer than the FS receiver's own band edge, to (conservatively) guarantee that the offset frequency AP's energy footprint within the FS link band is at least 11 dB lower (36dBm – 20 = 16 dBm) than that of a corresponding LPI footprint centered precisely on the FS link's center frequency. (As noted in prior sections, such is ubiquitously permitted without operational sanction anywhere in CONUS). So, a 1 MHz guard-band on either side of the FS channel to define a keep out region for overlaid AP energy is certainly sufficient. This has the takeaway specification that the nearest AP center frequency (Fc) for a 160 MHz channel can be no closer (on either band-edge side of the legacy FS link) than 81 MHz away to guarantee the necessary suppression. This can be shown pictorially as follows:

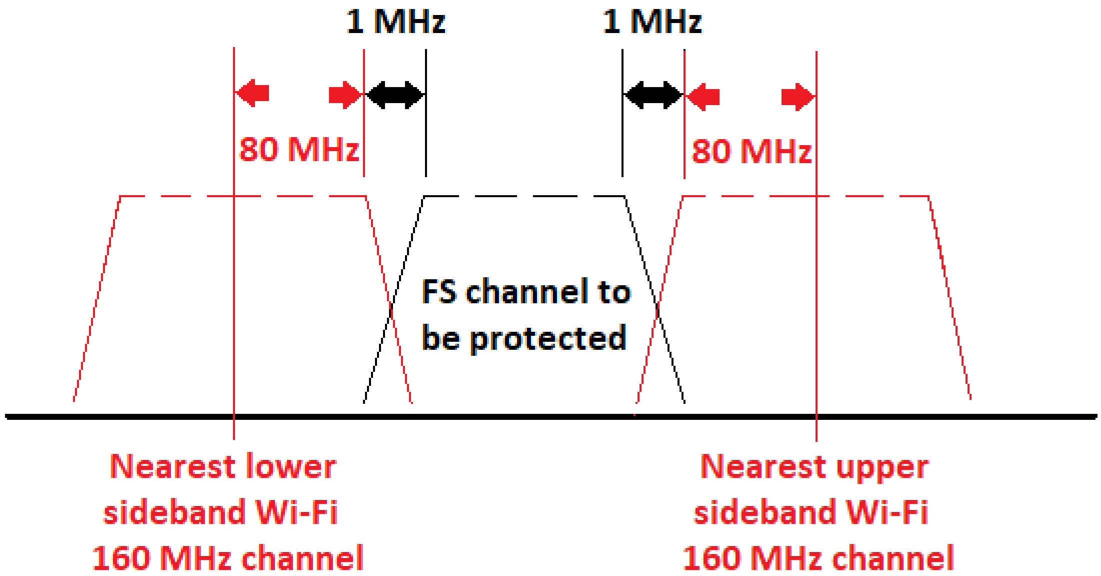

Figure 23 - Suggested FDM Keepout for Adjacent Wi-Fi Channels

In a spectral asset pinch (and as a last resort), Wi-Fi channel BWs less than the full 160 MHz may be considered for "spectrum puncturing" opportunity.

5.3.5 Ancillary AFC Services

The operational aspects of the AFC system in the prior sections have dealt exclusively with AP and cloud considerations for manufacturers and service providers. However, there is a consumer-facing aspect which needs to be acknowledged and dealt with: absent some type of *a priori* analysis, how could a consumer be assured that they would be able to leverage the standard power AP they are considering for purchase? The obvious answer seems to be that, like an AP, they should have the ability to query the AFC system (with appropriate geolocation detail) and have it deliver a verdict on the usability of a standard power device at the location they specify for use. Simply put, if they would not be permitted to exercise the standard power option where they reside, there would be no point in investing in a more powerful (and expensive) AP than one which exploits LPI.

The AFC specification also permits what amounts to a private data channel (vendor-specific extensions) which could be leveraged (along with a log of prior AFC communications) for client applications to provide some monitoring capability of the AFC interactions with the AP; such would be useful for determining the cloud-to-ground behavior of the AFC system and perhaps also assist with fault analysis.

5.4 AP Behavior Under AFC Management

The touchstone for managed AFC operation (and prerequisite for AP certification) is the compliance of the AP to the operating transmit power mask delivered to it from the cloud portal. That is, absent AFC connection and mask delivery, the AP cannot presume to operate at greater EIRP than that allowed for LPI operation. Further, it must default to no more than the LPI specification at lease timeout (currently specified to be 24 hours); finally, if granted a mask (as either available channels or frequency spans), it will not presume to a standard power EIRP in those spectral regions specifically masked off by the AFC query response and only leverage the spectrum allocated. In all cases, however, the AP will adopt a "listen before talk" spectrum deference posture (to mitigate interference from alternate, non-Wi-Fi, unlicensed links). This layered approach to spectrum exercise elicits several algorithmic methods to tune AP behavior for compliant operation and are examined below.

The following segments detail potential AP operating behaviors above the foundation presumptions just detailed above. Note that references to other bands besides 6 GHz are included only for completeness and acknowledge that 6 GHz band assignment for clients necessarily must account for the impact of mapping traffic to clients across *all* available Wi-Fi bands serviced by the AP.

5.4.1 Wi-Fi Airtime Engine (WAE)

The AFC-bound APs seeking to leverage standard power may comprise any of several multiband Wi-Fi options (up to, and including what some tag as "quad-band" APs – 2.4 GHz, 5 GHz low band, 5 GHz high band and 6 GHz). Assignment of multiband clients to 6 GHz (which define the magnitude of the exploit of 6E standard power) will require assessment by the AP regarding allocation of its band resources based upon projected bitrate consumption and the parametric contribution of several considerations. Among these are airtime budget/band available and in-use, desired margins to known airtime limits, band-use and bitrate-use profiles of the prospective client population, status of per-band data queues in the AP, limitations of the AFC-derived channel power mask, historical time-of-day prior behavior and detected desired-channel interference from other Wi-Fi APs or unlicensed 5G/LTE transmitters. The core determination effected is an assessment of airtime burn in each of the bands (especially 6E, given its higher degree of scheduling determinism and ultimate rate/latency performance). This evaluation we will refer to as the WAE.

As previously alluded to, consumed airtime per band is a result of supported bitrate at a particular set of MCS values; these latter are sustained by good C/N in the bandwidth being used – which itself is a function of radiated power and link length. The greater the EIRP and the shorter the path length (loss), the higher the potential MCS – and the less the airtime used to transmit the data packets. The reverse is obviously true: lowering the EIRP or extending the service radius to a far-flung client will lower the delivered signal C/N at the endpoint, reduce the operating MCS for the link and increase the airtime consumed to deliver the data packets (associated with that client, but nonetheless deducted from the overall airtime in the standard case of a single channel/band). Note that these aspects define the viewpoint from the AP to clients or peers (downstream traffic). Because the links support duplex operation, the corresponding upstream MCS and packet sizes must also be accounted for in the channel airtime usage. (And the upstream MCS will be lower than downstream in the 6 GHz band – given the Wi-Fi 6 FCC-mandated 6 dB backoff in client EIRP relative to the AP). The goal of the WAE is to canvas the available operating parameters and balance Wi-Fi band use to achieve best overall duplex data throughput (up to the SLA limits of the WAN).

The following diagram describes the algebra around estimations of band capacity (presume single channel/band) and band availability for future service mounts:

Definitions and Calculation

Per-band observations:

*** Duplex link bitrate capacity (C_L) defined by MCS_L, number of spatial streams and channel BW -- in both US and DS.**

*** Spatial stream count is the lesser of SS on either side of the link (typically client-limited).**

*** Examine all links and determine link bitrate capacity for each (historical or actual MCS, # of SS and BW). Discount the capacity ~ 20% to account for framing and TCP overheads.**

*** Examine all links and determine link bitrate demand (average packet size x average arrival rate -- pps).**

*** Calculate air time as the ratio of bitrate demand to bitrate capacity per link (T = D/C).**

*** Channel utilization =** $\sum_{\text{all links}} T_L$.

*** Channel availability = 1 - Channel utilization.**

Figure 24 - WAE Term Definitions and Airtime Calculus

Note that the calculation involves acknowledgment of some amount of Wi-Fi framing overhead, coupled with TCP loop closure (as roughly 10% of available goodput airtime). This allows us to estimate goodput bitrate limits when calculating MCS-based link spectral density (as bps/Hz of channel BW) and then have the framing overhead count against that calculated link capacity. Using these calculations per band allows the AP to weigh service mounts based on actual available capacity (versus merely physical layer (PHY) expectations) per supported Wi-Fi band.

The general mechanism of consulting cached (historical) service mounts on a per-client basis would allow the WAE to do an initial assignment of client device to band based upon detection of the client's presence in the network, said client's own supported bands (and MACs) and a review of the client's recent past data consumption. The key historical factors are observed link upstream (US) and downstream (DS) MCS and data rate (calculated as a rolling average of packet size and arrival interval) historically associated with the client(s) in question. Actual MCS detail will be sounded out when the client attaches to its assigned band and may differ from historical indications for the case of a different applications environment, relocated mobile client devices or assignment to a different band (whether the device is mobile or fixed). By invoking historical record, the WAE can make a predictive estimate of the band availability impacts when a particular client lights up on the network.

Two diagrams immediately follow, one dealing with the static interplay of entities or parameters which influence the WAE and a second to demonstrate band assignment strategies for given mounts of clients/services. The third consideration (accessing 6E standard power against a stored mask of available power/channel in the 6 GHz band) will be detailed after these two.

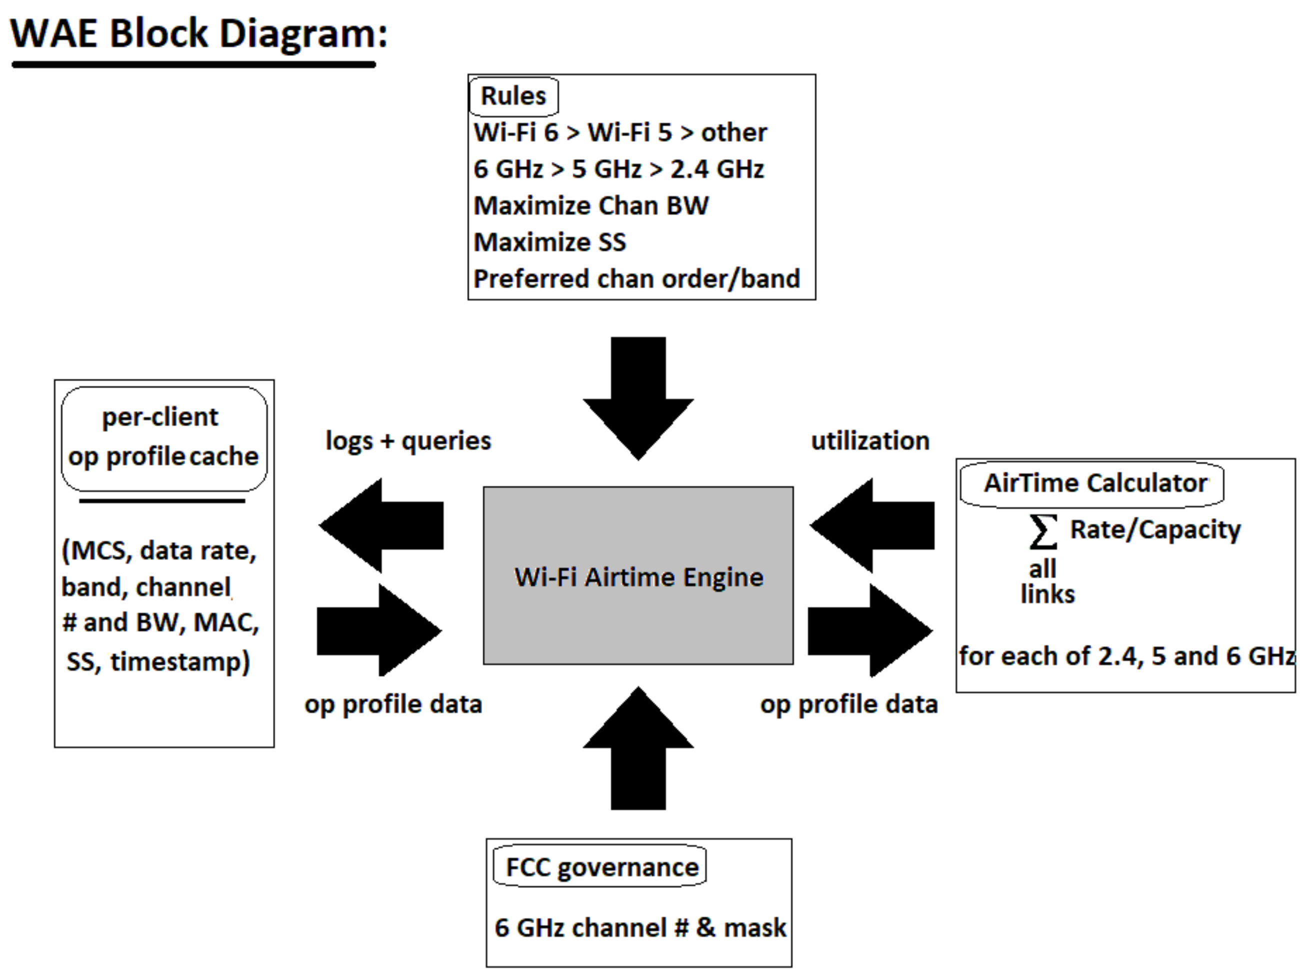

Figure 25 - WAE Interactions

The block diagram lays out the dependencies and interactions of the WAE and four operational data repositories: the profile cache, FCC governance, airtime calculator and client link parameter assignment rules. The rules block is essentially provisioned into the device in one of several ways: via cloud-based remote radio management (RRM), core software (and upgrades) or application-solicited user preferences and directives. The nature of the rules is to establish an order precedence to be used in the WAE's assignment of client operating parameters. Examples of these are shown in the diagram above; Wi-Fi 6 supported connectivity has a higher precedence than Wi-Fi 5 (especially at 6 GHz, where MAC homogeneity guarantees excellent, ungroomed, latency performance); which itself is the intermediate choice before indulgence of airtime-robbing older legacy Wi-Fi MACs. In similar fashion, where client band operation offers options, the preference would be to assign clients to 6 GHz before 5 GHz – with the 2.4 GHz band being reserved for devices which can only operate in that band. Generalized rules would be captured by notions such as maximizing channel BW and SS for any given link (to invoke the best bitrate – and hence, shortened airtime to deliver packets – for each client link). Other rules might establish channel number assignment in each band (based on prior attachment successes, say).

As regards the operating parameter cache, this defines a collection of most-recent client link parameter assignments over some finite rolling period with tagged timestamps. The cache reflects productive historical assignments which the WAE can reference in assigning initial operating points for a client link ahead of calculated airtime ratios based on the anticipated demands of any new service session. (This permits profiling of clients in terms of their data ingestion and production stochastics.) This should prove hugely beneficial for client devices which are single-purpose network mounts and unlikely to exhibit a wide spread of operating demands.

Finally, the airtime calculator maintains a monitor on the sum of all client link ratiometric (as rate/capacity) performance. Typical margin determination would assign a percentage (in the region of 20-25 %) of airtime excess in each band and the airtime calculator would establish band congestion based upon the available excess airtime available per band after the margin and all calculated airtime ratios are accounted for. This provides the WAE with an airtime budget per band which it utilizes in determining optimal client-to-band assignments (transferring bitrate loads as necessary to keep all bands at reasonable utilization levels).

5.4.2 Client Mounting Process

It follows then that these data repositories defined above can be leveraged in a controlled order to effect client band assignment for the AP WAE. An illustration of this follows:

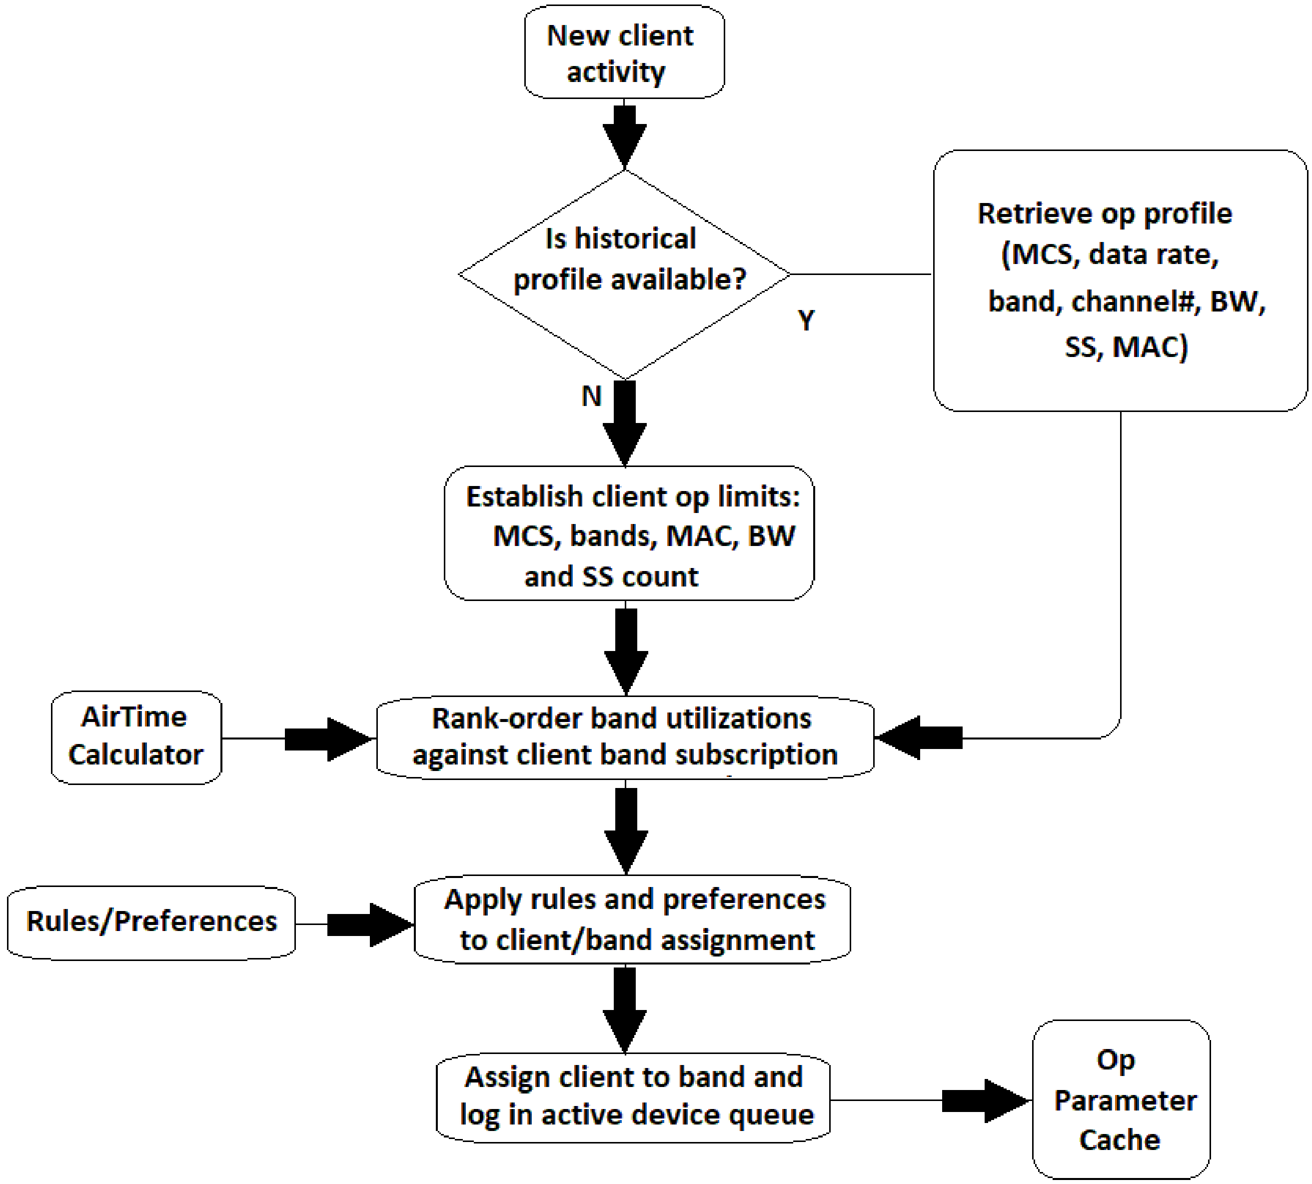

Figure 26 - Wireless Client Mounting Process

Note that housekeeping for the Op Parameter Cache is not included here but can be assumed to include retirement of latent client parameter detail for situations where the timestamps of the data exceed some "hold" timeout period.

The AP's decisions on band use and power/channel assignment are associated with time-of-day historical set points and available empty (or low CCI) channels. As regards the former, the WAE would make use of the Op Parameter logs and interpolate between those timestamped captures and the present time-of-day (TOD) to determine a matching set of parameters (modulo-24 hour).

5.4.3 The Importance of Channel Scanning

In the creation of an available channel map at 6 GHz the AFC-provided channel mask is the starting point of the process. However, responsible stewardship of the channels provided also requires the AP to develop a view of potential unlicensed competitors for the Wi-Fi channels granted it. (There is the unintended consequence, in multiple dwelling unit (MDU) situations, that nearby Wi-Fi customers may all receive the same AFC mask, which unfortunately serves to increase the risk of CCI on the available channels). RRM can assume the mantle of direction of the spectrum decisions; however, the process offered for consideration here is for the cases where RRM is either not available for the AP or there exist disparate, non-aligned RRM orchestrations proximate to the AP. The implication for the WAE is that it must maintain a background scanning function to determine a rank-order of acceptable channels to use (based upon detection of energy during the scan). The scan may be optimized for efficiency (time required to distill the exploitable spectrum) versus "freshness" of the available data. One such optimization would be a progressive triage of available channels based upon initial scanning at maximal bandwidth for the AFC-granted channels (from the collection of all but the current channel in-use) and then progressively narrower scans (by factors of 2, down to a limit of 20 MHz) of those channels which exhibit energy below some set threshold. The outcome could be used to create a candidate channel ranking down to an ultimate resolution of 20 MHz. (or to that bandwidth which has been determined – or commanded – to be the minimally acceptable one. It is certainly possible that, due to bitrate demands at 6 GHz, only options at 80 and 160 MHz may be considered). The algorithm (for determining the best option for a new channel in a single band) is logically illustrated as follows:

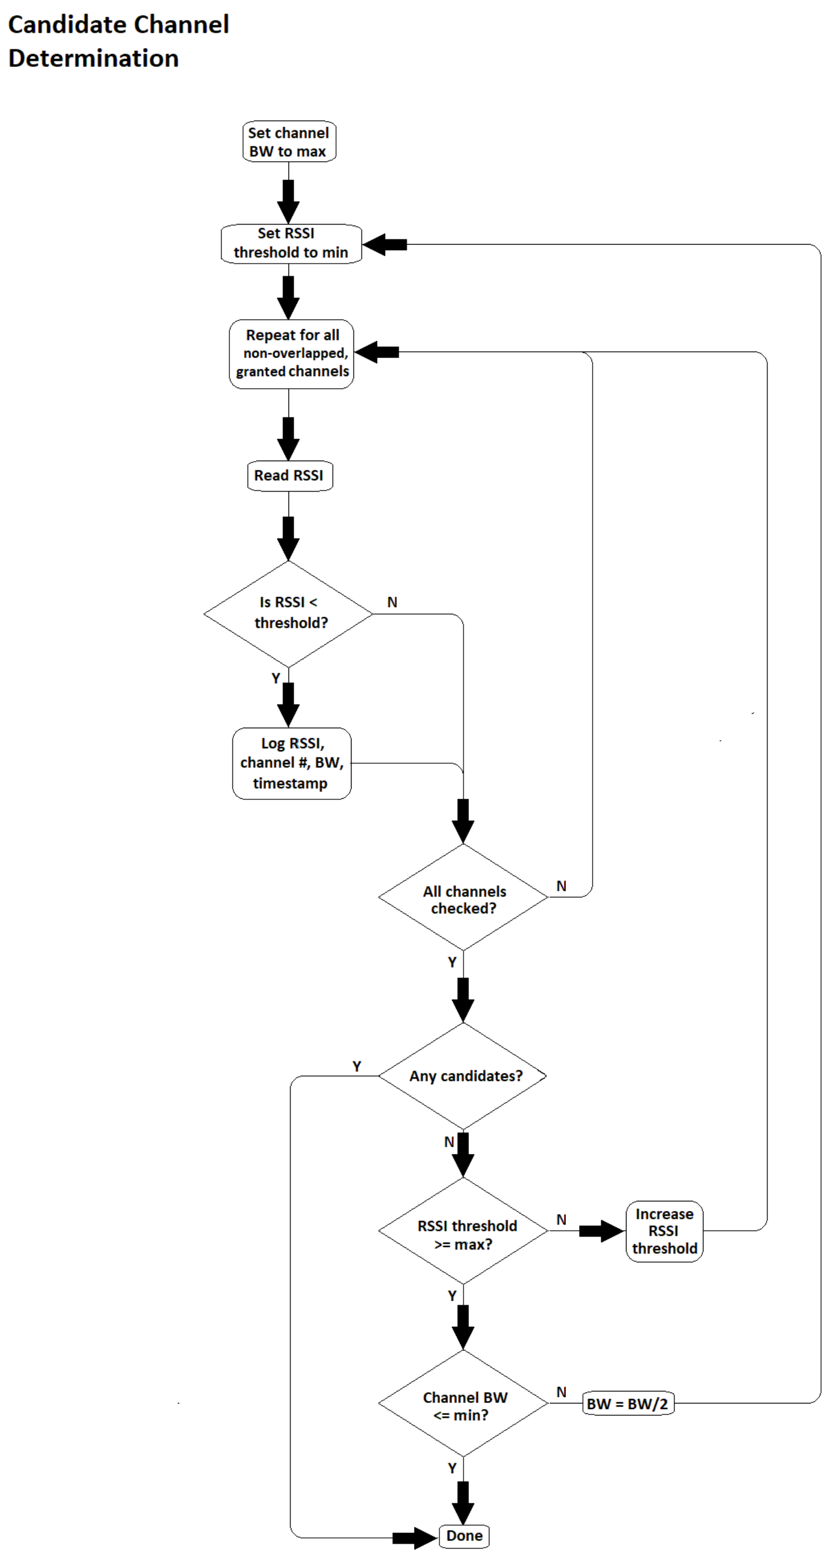

Figure 27 - Scanning Algorithm for Determination of Unconstested 6E Wi-Fi Channels

Given the still-pending maturity of the AFC specification, some of the details above are speculation – but in general, they represent the type of calculus and behavior a 4W AP would have to maintain to warrant its exploit of Wi-Fi 6E.

6 The Emerging Option of Indoor Adaptive Antenna Patterning

A brief aside on antenna developments appropriate to the various service mounts in Wi-Fi 6E is appropriate. Fundamentally, there are two classes of antenna arrangement which serve the interest of robust indoor Wi-Fi distribution: isotropic and narrow-beam patterns. As regards the former, it is expected that both client and AP groups will continue to support multi-direction antenna footprints for major slices of Wi-Fi airtime (to account for random AP-to-client spatial distribution and promote the least expensive method to implement fixed, full azimuth and elevation reach). Narrow-beam near-line-of-sight (nLOS) radiation patterns are, however, useful for P2P home inter-mesh trunk hauls, given that these endpoints are spatially fixed relative to each other. Their directional pattern distributes the allowed EIRP across a reduced spatial cone (improving link margin to one preferred, fixed receiver at the expense of sensitivity for multiple others) and thereby limits the scattering which promotes multipath-based symbol spreading. This unwanted aspect of isotropic radiation can dictate use of wider timing guard band intervals (for expansive floorplans) which can rob goodput throughput. It also inflates error vector magnitude (EVM) at a particular link operating point and hence, represents compromise of the delivered bitrate through potentially unwarranted MCS reduction.

In the scenario where bookended APs are used to establish a trunk haul between meshes (and involves at least a moderate concentration of 6E clients), it makes sense to consider smart antenna switching which, on a per client basis, alternates between isotropic radiation at some inclusive EIRP (which does not tax client receiver dynamic ranges for the respective -- likely 2x2 or even 1x1 -- clients of each AP endpoint), and a 4W, directionally specific, full spatial-stream-exploiting trunk link. The combined efficiency of sharing single-channel airtime between the two modes implies that a second channel (and radio) need not necessarily be resorted to; it is, at minimum, an option to consider for a high-end class of standard power AP.

7 Service Mount Considerations for Standard Power Operation

With the available recourse of higher EIRP, considerations of how best to integrate such a capability within a home Wi-Fi fabric can be evaluated. More transmit energy begets much-improved home wireless coverage, but there are architectural nuances to be evaluated. In addition to the increased bitrate comes the assurance of much better latency performance (between 10x and 100x, versus 5 GHz) for 6E links, primarily due to the single Wi-Fi media access control (MAC) to which every band participant must exclusively subscribe. At standard power, it should be possible to link large client populations throughout a spacious multi-story floorplan and still meet < 2 msec latency for all client links.

7.1 Client Backoff

By far the biggest impact on in-home Wi-Fi service throw at 6 GHz involves the FCC-mandated 6 dB backoff in client EIRP versus its servicing AP (whether this device operates at LPI or 4W). And it is at the point where client transmissions – even down to MCS0 – cannot be recovered by the AP that Wi-Fi framing collapses and the client is disenfranchised by the mesh. At this point in uplink communications, the downlink is still capable of decent performance – so the argument goes that downlink-biased services (most are of this type; the exceptions being media backhaul for home security or work from home – WFH – scenarios) will not sputter and burn airtime on retries but more abruptly, simply (and prematurely, perhaps, from the transmitting AP's standpoint) fail.

The implication to whole home LPI wireless service is clear: once a client outreaches the service throw, the user has no other choice than to insert a wireless extender into the mix. But as pointed out earlier, at standard power operation, these service radii easily accommodate very large floorplans even when the single AP is not optimally located (i.e., in a lower corner of a multi-story home, for example). So perhaps the better way of viewing standard power is that it affords the opportunity (up to well above American average home floorplan size) to obtain the desired WAN-matching service-level agreement (SLA) with only a single AP whose location is determined solely by convenient access to its WAN. For larger homes, a summary takeaway might be that one could guarantee > 1 Gbps client service everywhere throughout the home, even with a corner basement mounted wide-area network (WAN) gateway.

7.2 Subordinate APs

It is noteworthy that the FCC has created an exception category to the backoff for what it deems "subordinate APs" – essentially peer nodes on trunk links who themselves have one or more clients to service. For these devices (think "extender" in the case of a large home) EIRP need not be reduced. This has the effect of providing for exceptionally robust connections between mesh nodes simultaneously serving proximate battery-constrained clients and implies that the insertion of an intermediate hub is almost always guaranteed to enfranchise clients whose path loss to a single, non-optimally placed, AP renders them otherwise "unreachable". In doing so, it also appreciably raises the average link MCS throughout the whole home, which pares the overall airtime spend and makes single-channel radio extension at least conceivable.

Note that this exception to EIRP backoff establishes standard power Wi-Fi as a less costly option for large homes. Simply put, even if the standard power AP is a bit more expensive than an LPI-constrained one, it costs less (in dollars and APs) to cover a whole home to a given aggregate bitrate service level than to attempt multiple LPI extensions to keep the wandering low-power clients happy. And not to put too fine a point on it, fewer (or no) linked extenders imply a much more robust network availability factor.

7.3 Standard Power Deployment Opportunities

So now we come to advised exploit of 4W indoors for Wi-Fi 6E APs. Certainly, the span of client capabilities (and service bitrate expectations) will define the most efficient leverage of multi-power Wi-Fi meshes for given service mounts. Using the testing at LPI and the projected improvement in reach afforded by standard power, we can imagine several architectures which would efficiently serve even the largest home layouts and unlock advantages of 6E. We already have data which indicates mobile client support anywhere in the Wi-Fi house based upon the servicing LPI AP being located at the home midpoint; when this AP is promoted to standard power, the rates cannot help but improve (up to the uplink starvation case previously mentioned). But this turns out to be a needless worry.

Some examples of the expectation for client devices in the Wi-Fi house if the central AP is bumped to 4W would clarify the monstrous benefit of a 4W midpoint AP in the 5300 square foot home. Referring to our data, a repainting of obtainable TCP bitrates for 2x2 mobile devices scattered throughout the Wi-Fi house becomes a monotonous affair. No room on any floor is served with less than at least 1.4 Gbps downlink bitrate – and the entirety of the main floor (where resides the AP) sees the maximum 1780 Mbps. The tiling looks as follows:

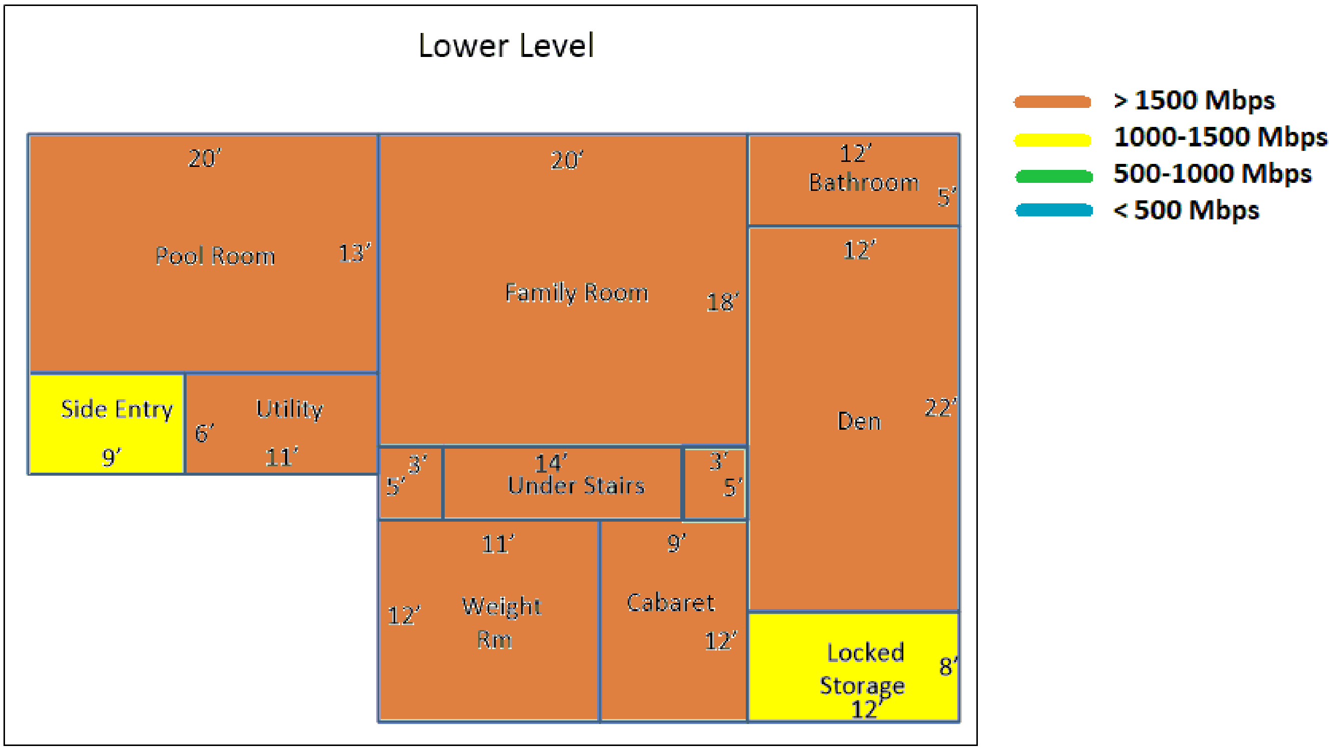

Figure 28 - Lower Level Service Coverage with 4W AP

In fact, though clients at the side entry and locked storage locations on the lower floor would be constrained to just under 1.5 Gpbs with an AP association from the home midpoint, in the case of the side entry location clients would be steered to a binding to the WAN gateway. Under this presumption the full 1780 Mbps becomes available there. Signal throw to the locked storage area, however, is actually worse to the WAN endpoint than the living room AP, so the modest attenuation of bitrate to 1.4 Gbps would remain.

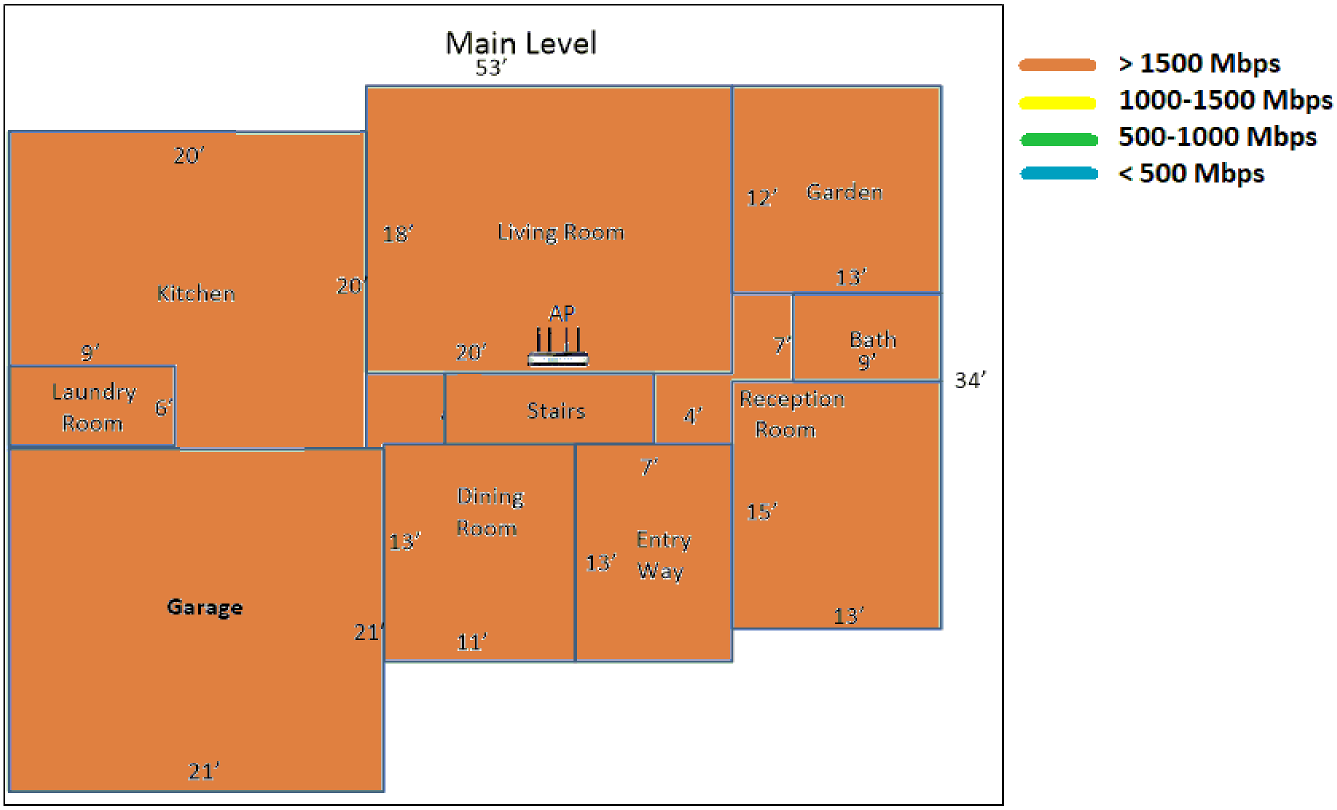

Figure 29 - Main Level Service Coverage with 4W AP

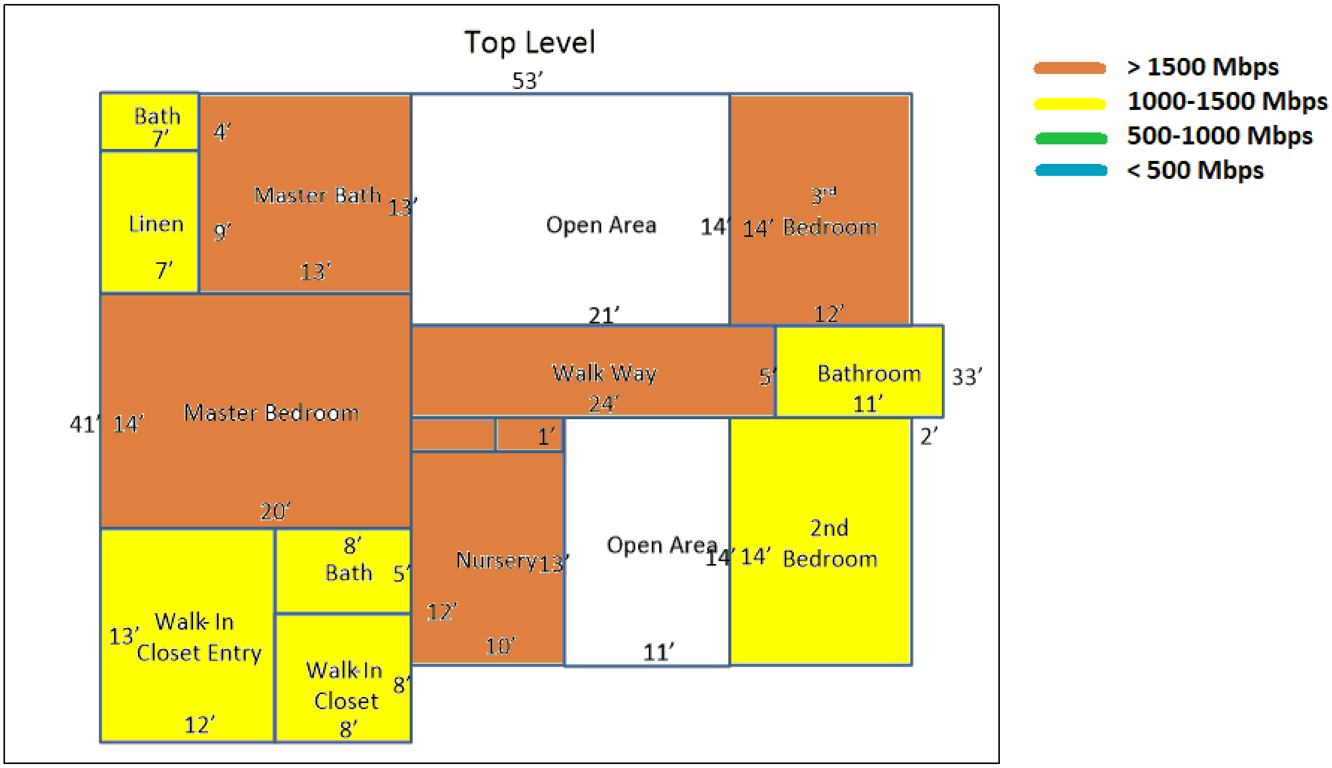

Figure 30 - Top Level Service Coverage with 4W AP

Noteworthy in all of this downlink behavior analysis is the comforting realization that, despite the rather onerous limitation of 20 dBm of transmitter power, the roaming mobile clients never experience uplink service anywhere in the house to < 550 Mbps (so the prior bookmark on potential uplink starvation may be retired).

But the midpoint approximation – not a bad alternative and easy enough to implement -- is likely not the best coverage option for extremely large floorplans. As the WAN attachment is in the basement pool room (a common enough home topology for connecting an outside WAN), it would make sense to move the extender all the way to the upper floor and closer to the two Jack/Jill bedrooms. This has the airtime benefit that more clients can be steered to the WAN-attached gateway, such steering subtracting airtime from the 4x4 trunk (and eliminating one link latency/client for those which can so attach). These latter outcomes are most welcome, in that the movement of the extender AP endpoint further away from the WAN will (slightly) penalize the trunk bitrate performance (so it having to carry less backhaul/fronthaul traffic tends to somewhat balance the capacity drop). The modified layout looks as follows:

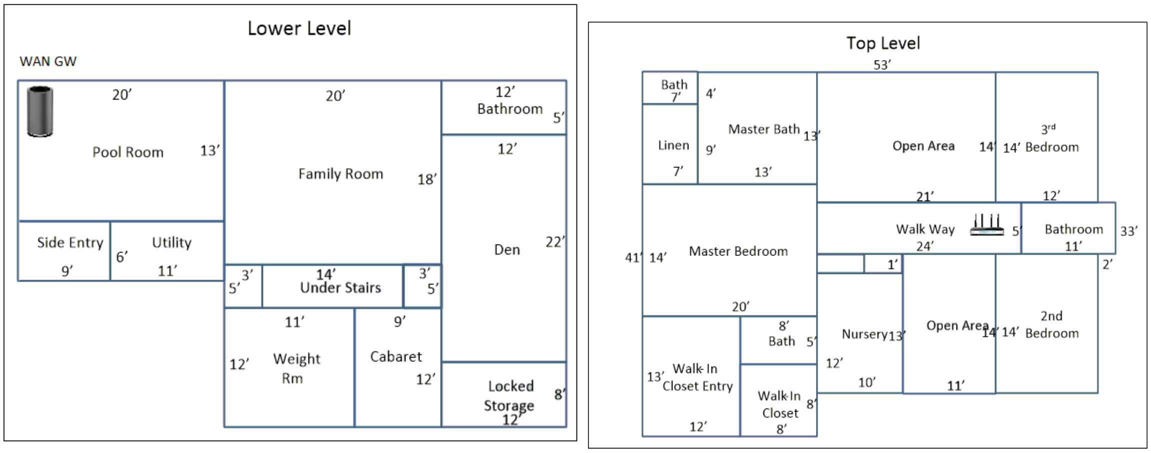

Figure 31 - Modified Trunk Topology to Improve Whole Home Coverage

In this scenario, trunk performance is somewhat degraded (recall the midpoint extender placement allowed for a 3.5 Gbps goodput carry) to something more like 2.5 Gbps. However, the client coverage everywhere in the Wi-Fi house (except for that poor locked storage location in the basement, which achieves "only" 1.25 Gbps now) easily averages 1.7 Gbps. For the case of low numbers of 6E clients in a home (as will be the norm for the next ~5 years), this ability to roam (or set up a fixed client endpoint) anywhere in the home and not experience rate throttling manifests obvious value. Note that, in terms of present-day services, there are precious few application bundles which, even in simultaneous aggregate on one device, approach a 100 Mbps/client downlink service requirement – so having 1.5-2 Gbps at your disposal everywhere in the home is a bona fide lottery win.

This client "overindulgence", then, begs the question of whether a single 4W WAN gateway in the low corner of the house could blanket it with sufficient signal throw to enable excellent (or even more-than-adequate) service bitrate in a home the size of the Wi-Fi house. Without performing a full volume analysis, it is instructive to examine what the worst-case client downlink/uplink support appears to be, given the geometries in play. Selecting the far corner of the second bedroom on the top level as a checkpoint, the service radius amounts to ~ 62 feet from the WAN gateway (with two floors and perhaps 3 walls to penetrate – though the two large open areas facilitate multipath reach to this client placement). In this far corner of the house, downlink would be expected to be 1030 Mbps and uplink, ~ 275 Mbps – a fair margin to uplink starvation which would otherwise defeat persistent client attachment. The clear implication is that one would have > 1 Gbps everywhere in such a large home, even with a single, inconveniently located, AP. And let's be clear – the Wi-Fi test house represents a floorplan which is 2x the US national average for size of new construction, single dwelling units (SDUs).

It appears, then, that the recourse of two 4W APs is not obviously necessary in up to at *least* 5300 square feet of living space -- which means we would reserve that strategy for extremely large floorplans (realtor estimation of the mansion footprint, for example, sets an 8000 square foot expectation on qualifying space to earn that moniker). And it is not clear that the recourse of 2 standard power APs to blanket such a homestead would represent much of an investment concern for users well-heeled enough to own it.

Statistically speaking, given the predominant 6E client target capability of 2SS/device (which bounds the 160 MHz BW delivery of services to around 1.8 Gbps TCP), the overwhelming majority of US homes would *at most* require a single standard power wireless gateway – no additional wireless meshing required.

8 Conclusion

There is great value to be had in policing and permitting standard power operation in the highly scheduling-deterministic Wi-Fi 6 GHz band – and so extend the promise of this vast spectral asset which the FCC released last year. Though limited-power mobile clients represent a challenge to the brute force expectation that standard power APs should always be considered, it is nonetheless true that the advantages of the expanded coverage footprint over LPI suggest that a standard power "lynchpin" AP find representation in all floorplans more than perhaps 1500 square feet, to properly anticipate and lever the multi-Gbps WANs which are in the offing. Furthermore, floorplans in the 6000-10000 square foot range can make use of a single, very robust 4x4, 4W trunk (with bookended -- full 4W EIRP at each endpoint -- APs) to seamlessly connect meshes with client devices up to the periphery of the expansive floorplans (and minimize the peppering of LPI extenders everywhere throughout the footprint). The FCC and the greater industry at large have taken great precaution in defining and (soon enough) implementing a 6 GHz band coexistence scheme which simultaneously protects legacy microwave infrastructure investments and liberates all the promise of Wi-Fi 6E. The in-home high speed wireless future has never looked so bright.

Abbreviations

3D	three-dimensional
5G	Fifth generation of mobile communication technology
AP	access point
AC	alternating current
AFC	automated frequency coordination
BEL	building entry loss
CCI	co-channel interference
CGI	computer-generated imagery
C/N	carrier-to-noise ratio
CPE	consumer premise equipment
CONUS	continental United States
dBm	decibel milliwatts
DS	downstream
DUT	device under test
EIRP	effective isotropic radiated power
EVM	error vector magnitude
FCC	federal communications commission
FDM	frequency-division multiplexing
FEM	front-end module
FS	fixed service
FSPL	free-space path loss
FSS	fixed satellite service
FWA	fixed wireless access
Gbps	gigabit per second
GEO	geostationary equatorial orbit
GHz	gigahertz

I/N	Interference-to-noise ratio
LOS	line-of-sight
LPI	low power indoor
LTE	long-term evolution
MAC	media access control (layer)
Mbps	megabits per second
MDU	multiple dwelling unit
MHz	megahertz
MS	mobile service
nLOS	near-line-of-sight
NRA	national regulatory authority
OOB	out-of-band
P2P	point-to-point
P2MP	point-to-multi-point
PHY	physical layer
PSD	power spectral density
SDU	single (family) dwelling unit
SLA	service-level agreement
SS	spatial stream(s)
STB	settop box
SUT	system under test
TCP	transmission control protocol
TOD	time-of-day
UDP	user datagram protocol
ULS	universal licensing system
U-NII	unlicensed national information infrastructure
US	upstream
WAE	Wi-Fi airtime engine
WAN	wide-area network

Bibliography & References

Frequency Sharing for Radio Local Area Networks in the 6 GHz Band, ver 3, RKF Engineering Solutions, LLC, January 2018; prepared for 6USC

ET Docket No. 18–295 and GN Docket No.17–183; FCC 20–51; FRS 16729, *Unlicensed Use of the 6 GHz Band,* FCC

COMPATIBILITY STUDIES IN THE BAND 5725 – 5875MHz BETWEEN FIXED WIRELESS ACCESS (FWA) SYSTEMS AND OTHER SYSTEMS*, 2005, ECC Report 68*

