Operational Sustainability: On Achieving Optimal Leverage of the Power Grid with Wi-Fi CPE

A Technical Paper prepared for SCTE by

James R. Flesch
Director Adv. Tech, office of CTO/Home
Commscope
3871 Lakefield Drive, Suwanee, GA 30024
jr.flesch@commscope.com

Brian Carroll
Sr. Manager, HW Eng'g/R&D
Commscope
101 Tournament Drive, Horsham, PA 19044
215 323-1427
brian.carroll@commscope.com

Table of Contents

List of Figures

1. Introduction

Sustainability has established itself as a critical aspect of customer premise equipment (CPE) design goals, worthy of equal consideration to elements such as packaging aesthetics, thermal behavior, economic value proposition and overall electronic performance. While for some time this has drawn attention to items such as single-use plastic exposure, recyclable or otherwise non-hazardous material exploit in hardware implementation and some fair amount of device and accessory packaging cleverness, the interest now has expanded to include operational power signature which conforms to device service usage demands. Adapting power grid draw to match measured (or anticipated) application performance on a "just necessary" basis (and so eliminate wasteful standby consumption which – on a scaled basis – irresponsibly overtaxes the energy grid) is now the stuff of practical necessity. This type of considered and throttled behavior on the part of CPE immediately substitutes perceptions of thoughtless energy banditry with environmentally aware, green energy iconography -- which clever SPs can leverage in the marketing realm to amplify their brand's public association with responsible (and astute) corporate behavior. It also aligns electronic CPE with application-fitting power consumption as exhibited by other in-home electrical appliances – substituting managed operations for the manual ON/OFF cycling associated with their draws on the power grid.

There is a power consumption history associated with this class of CPE (Wi-Fi access points (APs) in particular) which must be acknowledged – an undeniable progression of unmanaged power sink which classically has not regulated itself during periods of constant feature aggregation and employment. As Wi-Fi medium access control layers (MACs) and bands have evolved and the number of spatial streams increased, so too has the power requirement to operate these devices. What used to cost approximately 5W for single band APs a dozen years ago has now risen to over 30 watts (W) for a tri-band concurrent (TBC) device with multiple internet of things (IoT) radios.

For this paper, reformulation of the power footprint for always ON Wi-Fi APs (be they wide area network (WAN)-attached or other) is examined to investigate approaches into making this class of device a better power dissipation partner in the home. We will examine tactics which lever hardware, firmware and software – the latter both device-native and cloud-hosted – with an eye to fitness for purpose (i.e., meeting the bitrate and latency requirements of services mounted at any particular time) and efficacy of power-saving results. The general form of solution – to appropriately hibernate various resources when service demand does not mandate their exploitation – will be a guiding rubric throughout this dissertation. Note that a collateral goal will be seamless transition from various hibernation states to appropriate service support levels which will eventually manifest as "slow to hibernate, quick to wake" behavior; the goal being such does not express itself in human-observable debits in local area network (LAN) performance, nor requires human intervention to achieve.

To enhance the end value proposition, the device will be outfitted with self-monitoring of the power consumed during operations and such will be made available to a cloud-based observer for both reporting purposes and as an aid in formulating pre-emptive behaviors based on prior history and time-of-day (ToD) considerations – the latter a key telltale input used to determine responsivity requirements of transitional device operations (to/from hibernation). Motivation for the implications of this operating litmus can be found in the following figure showing diurnal/nocturnal cable network operating bitrate demands (from March 2020) as can be inferred from ToD service group (SG) waxing and waning over the course of multiple 24-hour periods (and further grasping that such represents ensemble LAN back/front service hauls over that same period):

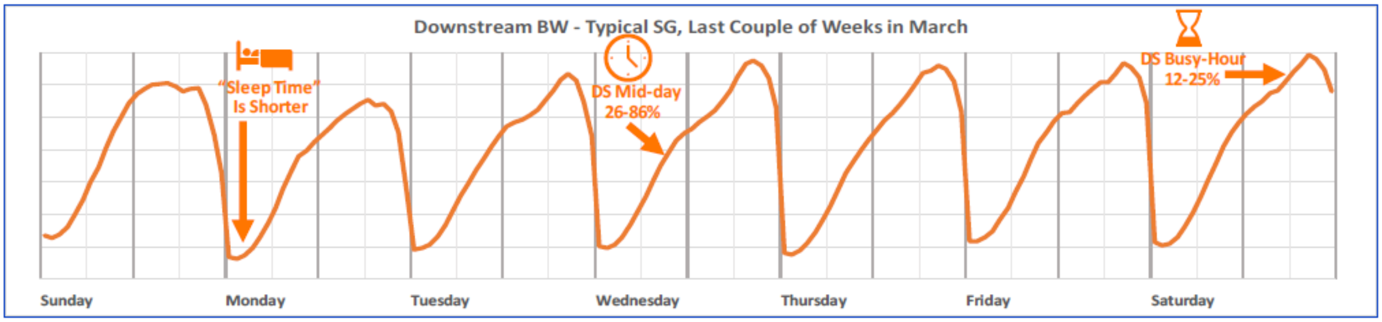

Figure 1 – WAN Bitrate Cyclic Variation Over One Week

Clearly, there is a ToD-based cyclic demand for wireless bitrate services which can form the master basis for modulation of power consumption on the part of a Wi-Fi AP.

These diurnal cycles of peak usage beg a simple brute-force solution – and it is fair game to point out the opportunity. As with all appliances, a simple ON/OFF strategy would serve to minimize power draw of the AP and, in cases where no keep-alive 24-hour services have it for a dependency, end up guiding us to the asymptote of power savings of which we should be mindful in targeting savings. For example, the expedience of 16 hours ON and 8 hours OFF buys us ~ 33% worth of power savings for the device and may be implemented by leverage of an explicit switch or a ToD-gating function:

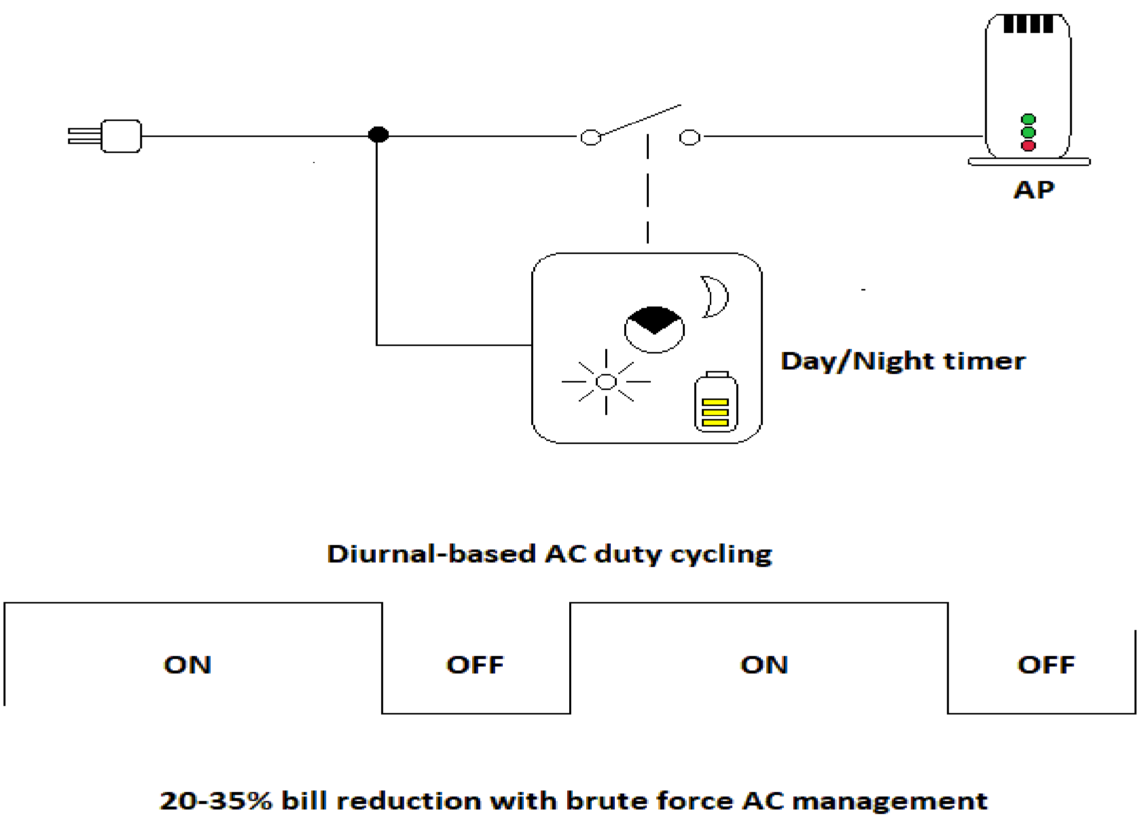

Figure 2 – Brute Force AC Management

Assuming, however, that AP availability for network attachment is associated with some support for 24-hour applications, we will proceed to break out design approaches which move across an operational spectrum from some minimally effective hibernation state to maximal performance. The European Code of Conduct (eCOC), for example, targets a bit more aggressive results in power use mitigation than that above – from mid-30W type of maximum dissipation to 8W or less in standby – so additional performance mining along this front should be conducted.

2. Scaling the Wi-Fi AP Hardware for Fractional Operations

In Wi-Fi APs, legacy design approaches have always proceeded on the basic calculation of how best to achieve maximum Wi-Fi signaling coverage for a given material cost point. Power supply specification, then, builds on presumptions of high duty cycle transmit leverage of all radio frequency (RF) spatial streams (as concurrent use), in concert with simultaneous support of all LAN and IoT interfaces. Central processing unit (CPU) core functionality is presumed to full extent of the available processing Dhrystone million instructions per second (DMIPs), and all corollary memory support is biased and clocked at maximal rates as well. Typically, this also invokes full use of active cooling (where such is provided) and so fan speed is also ratcheted to its upper limit. Against these artificially coincident maxima, packaging is then reduced to a three-dimensional (3D) envelope which balances best antenna farm orthogonality (as in- and cross-band isolations, typically the result of best physical separation of all coincidentally polarized elements) along with a maximally tolerable external package touchpoint temperature (or margin to silicon core temperatures, whichever is the more restrictive). Service provider (SP) architects routinely request small-as-possible industrial designs (IDs), so a balancing act occurs between a size minimization push and temperature/RF orthogonality pulls. There is a sort of comfortable challenge here, if it amounts to an incomplete accommodation for what we now wish to additionally accomplish: scaling the electrical draw of these maximal functions to dissipate appropriately fractional power loads when operational demands (service mounts) recede from these outer edge-of-performance considerations – thereby minimizing standby power use.

In this countervailing discipline, determination of required functional support from the AP must be pulled into a matrix which ends up defining which circuits/capability must be ON – and if frangible, to what order of inclusion -- for particular service profiles to be enabled. Bulk switchable aspects might include functional modules such as IoT radio support, battery charging circuits, Wi-Fi band(s), LAN ports, etc. Tunable aspects would include items such as number of RF spatial streams to enable per band and amount of CPU power (as clock speed based or other segmentation) and perhaps the amount of memory to apply to reduced operations. (This latter nut is a bit tougher to crack; responsivity in fractions of a second to some few seconds up to full service demand does not permit time for reconfiguration of the AP so unused software (SW) application containers in still-active random access memory (RAM) would remain an unavoidable – power consuming -- constant. (As circumstances have it, such memory hibernation consideration does not yield significant power savings in any event and so the complications associated with its consideration are rendered moot – refer to the next section.)

Establishing a budgetary feel for various circuits' impact to power dissipation is a necessary step. The following figure provides a weighting comparison of different functional elements to total power for a representative tri-band high end AP with 12 spatial streams (SS) (maximum dissipation around 35 Watts):

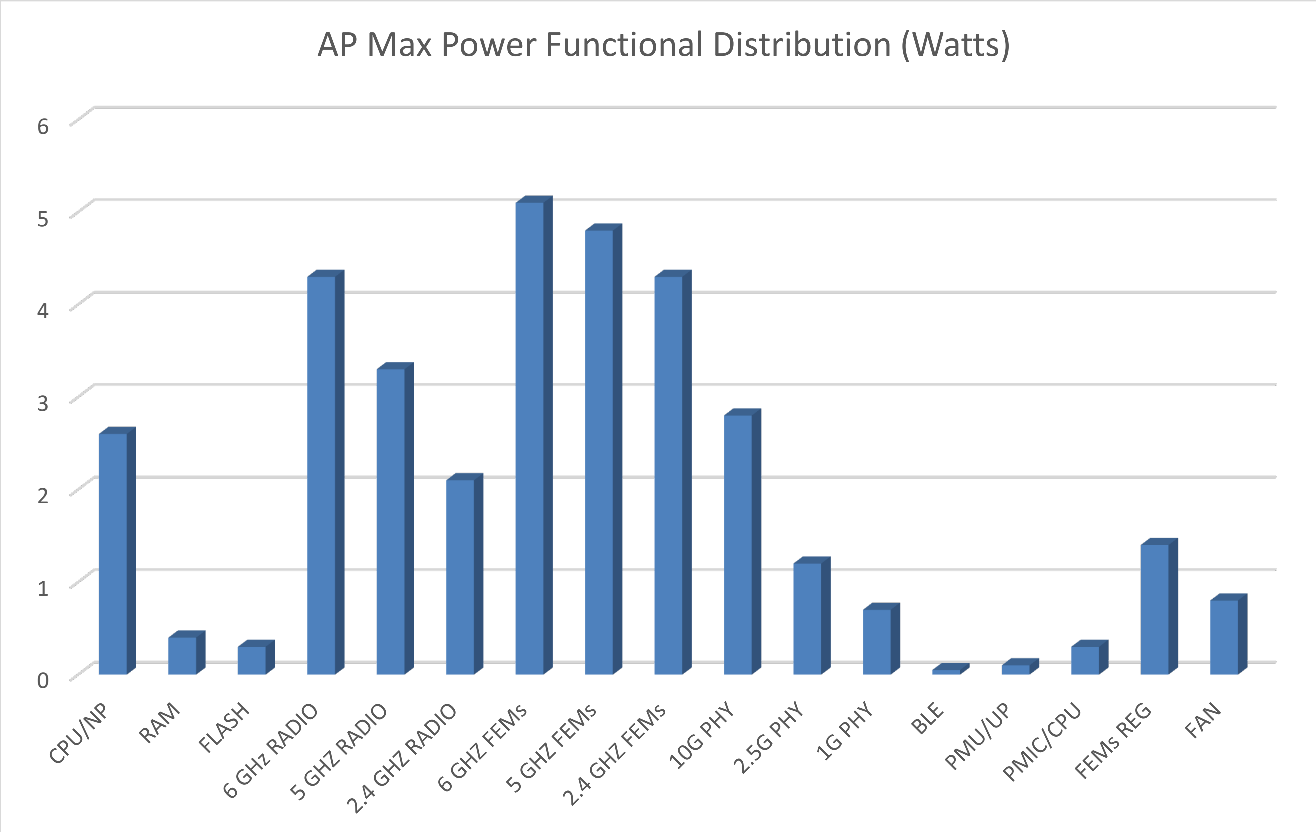

Figure 3 – Distribution of Maximum Power Consumed by TBC AP for Various Functional Blocks (~ 35 W total)

As can be seen, there are some broad power-management opportunities along the radio band, SS deployment, Ethernet ports and network processor fronts, with the majority benefit accruing to Wi-Fi radio management. These opportunities will be examined in more detail following.

2.1. Installing "Always ON" Test and Measurement

Recall the comment during introduction that self-reporting energy consumption will be a key performance aspect of the changes wrought to the AP. This metric is easily distilled, however, by simply monitoring voltage and current at the AP's power entry point (an external alternating current/direct current (AC/DC) brick form factor conversion being a common enough implementation). But the question of where to route the measured analog-to-digital converter (ADC) values and subsequent calculation (along with the hosting environment for this circuitry) needs to be answered as well; for the time being, we can assume a tagged sequence of calculated power values generated over some repetitive sampling interval is stored away for eventual routing to a cloud entity for analysis. How this measurement and storage is effected depends on the nature of the standby operations – whether these include IoT and wireline operations, for example, and if so, which physical layers (PHYs) are required to be supported.

Unfortunately, critical 24-monitoring typically accompanies medical and security services which straddle at least two different PHYs: Z-wave or Zigbee for security and Bluetooth Low Energy (BLE) for medical/wearable devices. (It is possible that in the very near future, Matter's Thread meshes will be able to see to both of these client communities, but available legacy devices currently occupy the different radio camps described above). If the IoT hub support is separately hosted, a very standard arrangement sees a single- or dual-band Wi-Fi backhaul or wireline Ethernet used to provide connectivity to the WAN. (Note that this is not guaranteed to be the case; arguably, WAN connectivity for IoT services might leverage direct-to-outside personal area network (PAN) – long range wide area (LoRa), for example -- or long term evolution/fifth generation (LTE/5G) committed cellular backhauls but for purposes of evaluating AP hibernation energy costs we will tread the path of presuming it proxies the ultimate WAN connection for IoT as well as 802.11-based radio services.) Given at least vestigial monitoring tasks during nocturnal operation (such being an appropriate time to idle – to the greatest extent possible – AP network processing and routing functions), the substitute of a cheap, very low power custodial processor to oversee AP hibernation and measure power use seems appropriate. The following figure details a possible approach to the very deepest level of hibernation, where nothing aside from IoT monitoring is presumed:

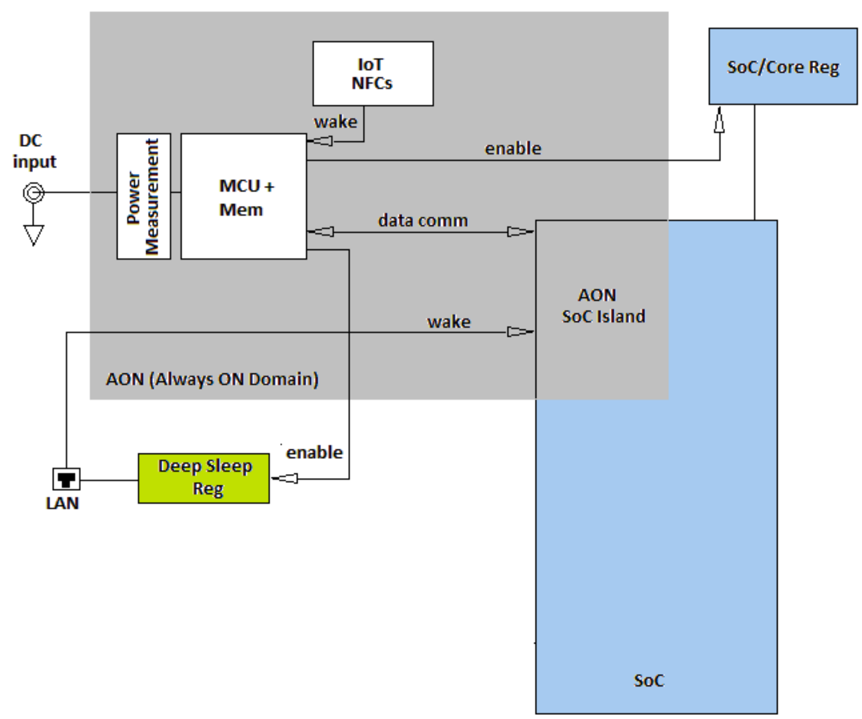

Figure 4 – Deep Hibernation of AP "Live" Sections With IoT-only Monitoring

In practice, it may also be a very good idea to permit some level of Wi-Fi awake triggering to occur, so the figure above would be modified to include minimal detection of at least the 2.4 gigahertz (GHz) Wi-Fi band in addition to the IoT and LAN emergency detection scheme. Note that from a WAN connectivity standpoint, the AP is effectively OFF at this point and so the timestamped power measurements would likely become bookended at OFF cycle start (power down) and OFF cycle end (just prior to power up). The decision on memory powering comes down to power up responsiveness. At an additional cost of ~ 0.7 W (0.4 W if only RAM is kept alive), the APs ability to light up with full containerized application space (versus a period of re-installing the software container environment) may be a worthwhile idle power investment.

2.2. Proportionate Network Processor Operations

During full or partial hibernation, the expectation is that executive management of the AP would see reduced activity from one or more of its multiple IP or other service flows and so the Network Processor/CPU could throttle back operations accordingly, maintaining sufficient awareness to react when service demands re-mount. This type of throttling typically involves clock speed reduction at a minimum (power dissipation being proportionate to clock speed) – but may also involve retirement of certain cellular functions as well, depending upon SoC power control granularity. The scalability is rooted in internet protocol (IP) packet traffic tending and so is proportionate to mounted service density; the goal is to invoke just enough DMIPS to see to the packet segmentation and re-assembly (SAR) and routing requirements (among other keep alive operations).

2.3. Partial Band Coverage

The halo product in AP space is considered a tri- or quad-band device (2.4 GHz / 5 GHz Lo / 5 GHz Hi / 6 GHz) with a MAC epoch of 6/6E (Wi-Fi 7 coming early next year). With currently available client devices spanning all these bands (rarely if ever all supported concurrently per client device, however), operation of the AP during peak hours could mean pressing all radio bands into service for a household with multiple devices. But clever "service aware" scheduling can counter this broad swath approach by examining the actual band service and interference profiles in a given household along with the quality of service (QoS) expectations of the various services -- herding capable clients onto as few radio bands as would meet the ensemble QoS proposition. (This herding is sometimes referred to as "band coalescing".) The operational behavior will be outlined in SW sections following, but from a hardware (HW) standpoint, the ability to isolate power dissipation on a per-band basis is key. This involves idling PHY and MAC operations on a per-band basis for those bands without active clients and assumes that device advertisement is best (first) serviced on the 2.4 GHz band – the expectation being that this band is supported on all clients regardless of Wi-Fi MAC epoch. For those rare cases where such client band support does not exist at 2.4 GHz, a minimal monitoring cadence can be maintained on 5 GHz as well. In "perfect hibernation", all client devices would be serviced at only 2.4 GHz and AP transmit opportunities (TXOPs) would be as infrequently subscribed as the minimal service keep-alive requirements permitted – and this, at a single active spatial stream and at a narrow bandwidth (BW) which still meets QoS demands for the near-dormant services.

Deviations from this asymptote would only occur for escalations of connectivity (alarm propagation of one type or another or the onset of service-heavy diurnal periods). Referring to the Figure above breaking out dissipation category, note both the granularity and impact to device power use based upon hibernating select band utilizations; avoiding 6 GHz, for example, in this standard power-capable TBC AP prunes roughly ¼ of the *total* power use from the budget by itself (if the band is fully de-activated). In the forthcoming couple of years, the adoption rate of 6 GHz-capable clients (coupled with the onset of post-Wi-Fi 5 MACs) will still lag significantly relative to the legacy band (and MACs) leverage, so the sacrifice as first resort seems advisable; the developing premises service profiles will indicate when (and how pre-emptively) 6 GHz use should occur.

2.4. Pruning Spatial Streams

Even in those cases where a particular band is drafted for use, circumspect employment of less than a full complement of spatial streams may still provide the necessary bitrate support while hibernating transmitters; however, this becomes a complex problem to undo – bitrate is inversely related to airtime, which is the variable (x SS) which defines the power profile for a given RF band. Remember too that bitrate is directly related to channel BW – increasing this to the maximum possible results in the shortest airtime transmission for a given packet length but consumes the most transmit power. For now, we defer the algorithmic impacts to the SW section.

In cases where internal frontend module (iFEM) (frontend modules (FEMs) internal to the radio SoC) technology is available, such simplifies the selection of spatial width to transmit over; however, it must be noted that these cases do not typically support power levels associated with fixed AP devices and are more intended for mobile, battery-powered devices. One further aspect is noteworthy: recent developments in FEM technology have seen improvements in idle (receive mode) operation, so full disabling of the external FEMs does not yield power savings during receive operations as it would have only a year or so ago (i.e., receive mode operations have become much more power efficient); nonetheless, given the inexpensive nature of pass device disabling of the bias for FEM blocks, it still behooves us to consider them for complete granular control of RF chain power usage (if the radio chip does not natively support individual idling of RF chains itself).

2.5. The Resulting Block Diagram

As a result of the considerations outlined in prior sections, we end up with the following AP ensemble block diagram which includes the granular controls required to selectively enable or disable the various contributing elements to total power dissipation:

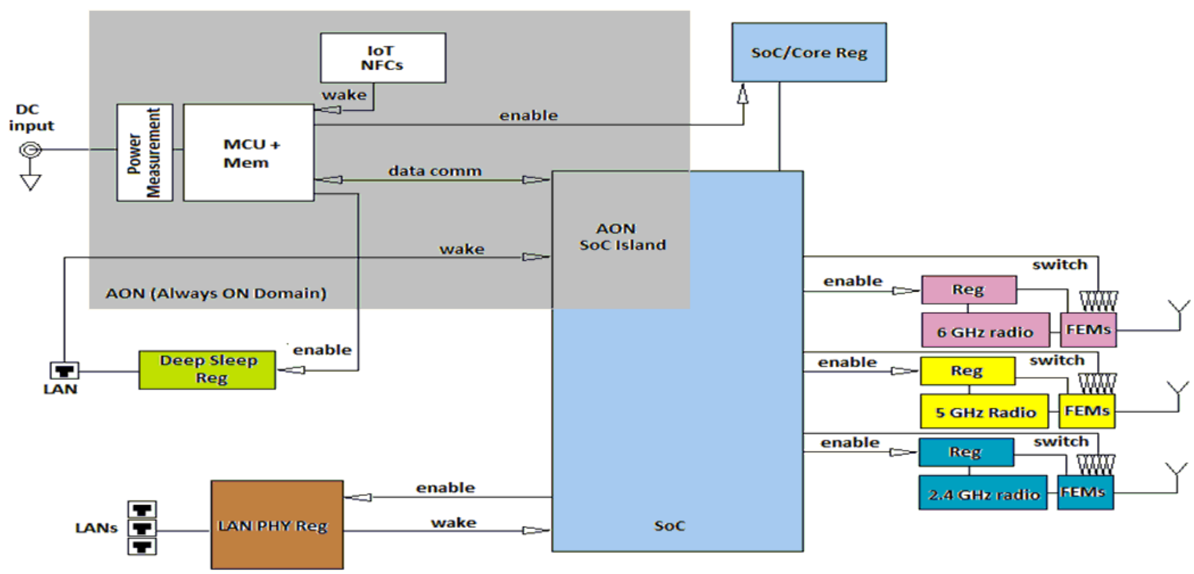

Figure 5 - AP Full Block Diagram with Granular Power Controls

2.6. Controllable Impacts to Power Budget

Selective enablement of the various control aspects detailed above result in a cascade of power operating points for the AP. The impact of major blocks is summarized below, along with the presumptions made for partial leverage of each of the subgroups:

Power Management Configurations						
	Off Mode	Deep Sleep	Network Standby	Active	On Full	**Conditions**
SoC	**N**	**10%**	**50%**	**75%**	**100%**	CPU Load
DDR4	**N**	**N**	**Y**	**Y**	**Y**	
eMMC	**N**	**N**	**Y**	**Y**	**Y**	
6G Wi-FI IC	**N**	**N**	**25%**	**50%**	**100%**	Duty Cycle
5G Wi-FI IC	**N**	**N**	**25%**	**50%**	**100%**	Duty Cycle
2.4G WiFI IC	**N**	**N**	**25%**	**50%**	**100%**	Duty Cycle
6GHz FEM	**N**	**N**	**1%**	**50%**	**85%**	Duty Cycle
5GHz FEM	**N**	**N**	**1%**	**50%**	**85%**	Duty Cycle
2.4G FEM	**N**	**N**	**1%**	**50%**	**85%**	Duty Cycle
10G WAN PHY	**N**	**100**	**100**	**1000**	**10G**	Port activity
2.5G LAN PHY	**N**	**100**	**100**	**1000**	**2.5G**	Port activity
1G LAN PHY	**0**	**0**	**0**	**1**	**2**	Ports Active
2.4G BLE IoT	**Y**	**Y**	**Y**	**Y**	**Y**	
PMU	**Y**	**Y**	**Y**	**Y**	**Y**	
PMIC	**N**	**N**	**Y**	**Y**	**Y**	
Losses	**N**	**N**	**Y**	**Y**	**Y**	
Fan	**N**	**N**	**25%**	**50%**	**100%**	
Power (W)	0.15	1.6	6.2	20	34.5	

Figure 6 – Dissipation Assumptions for Major AP Operational Modes

The categories shown include On Full, Active, Network Standby, Deep Sleep and Off Mode. On Full is a typical specification description which presumes worst-case concurrent operations and is used to investigate full circuit cooling and maximal Wi-Fi simultaneous use at high power levels. Ethernet componentry is also full exercised, as are all onboard regulators and the IoT subsystem. It is only a design bogie (and never encountered in normal operations), but nonetheless establishes the outer power envelope in which an unmanaged AP will inevitably waste power under partial loading.

Active Mode is a bit more circumspect in traffic handling – though it still tends towards the "over-exercised" portion of the operational spectrum. Wi-Fi use is still quite heavy and concurrent across all 3 bands, with roughly 2 gigabits per second (Gbps) packet traffic split across the two faster Ethernet ports. In "real" operational considerations, this represents a heavily loaded device (as might be the case during evening peak use cycles after 5 pm and before 11 pm or so). The power fingerprint by section looks as follows:

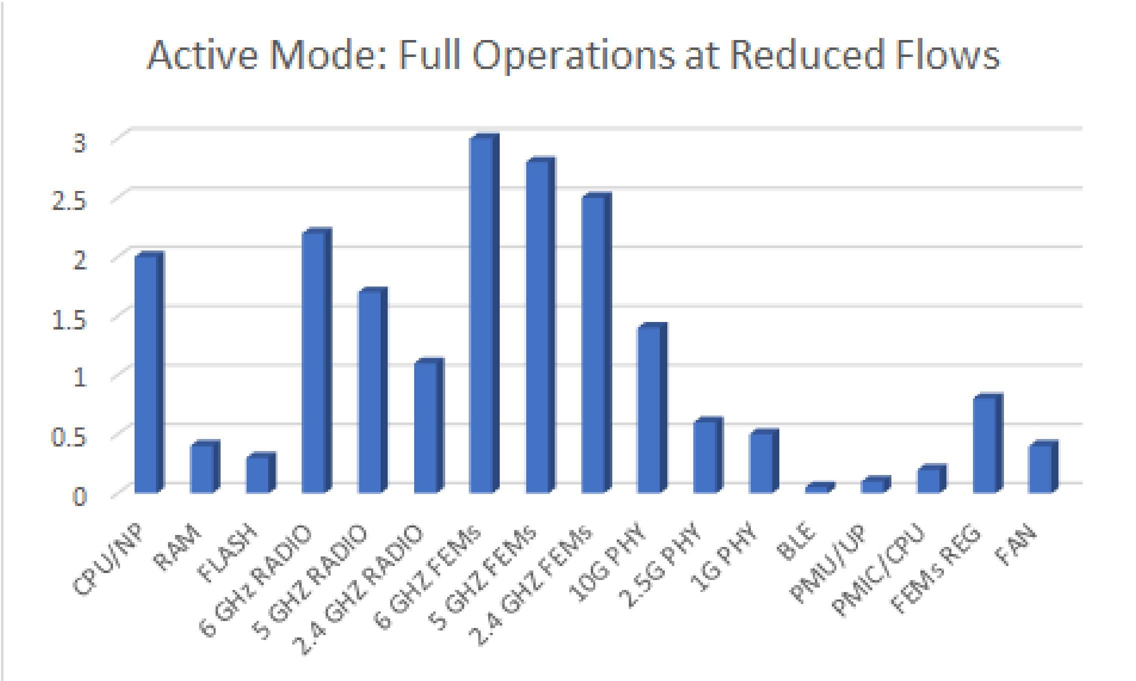

Figure 7 – 20W (>40% Reduction off of Max) Active Mode Power Distribution

Network Standby is, in declining power order, the next state shown in the table and essentially stands as the gateway between an operating AP and one of the two Sleep modes shown. Examining the table-supplied operational states definitions, we see that the device, while still active on all Wi-Fi bands is essentially just advertising presence and awaiting service mounts. At 6 W dissipation, it fundamentally establishes a lower bound on the most common operating range for the high-end AP: 6 – 20 W dissipation; it is this range we will attempt to shift lower by virtue of a slower cadence to awaken and lever the various Ethernet ports and Wi-Fi bands. Here is what comprises the contribution strata to this lower power level:

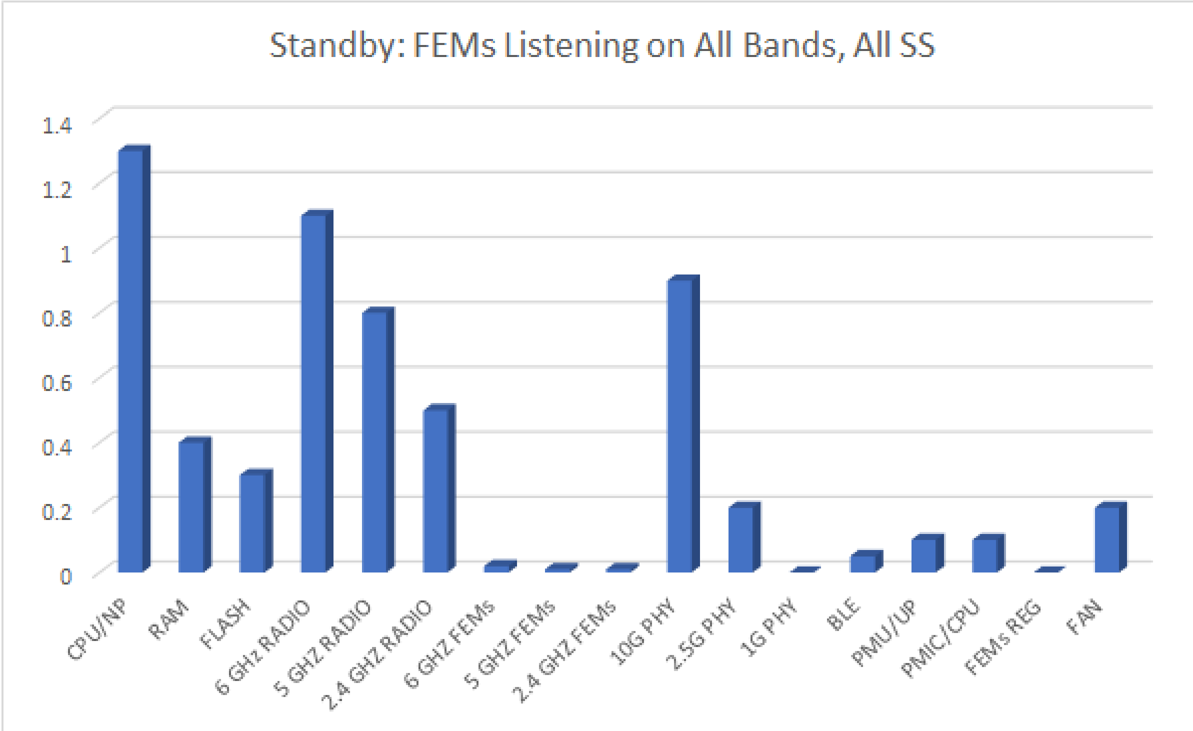

Figure 8 – Standby Power Profile, 6W

This brings us to the most dormant device states: Deep Sleep and "Off" (italicized in acknowledgment of the fact that some awareness is yet available). Regarding the former, the AP still has WAN connectivity; but to exploit this, memory may have to be energized in order to forward IoT-based alarms. Note that leaving the memory up (allowing immediate dispatch of same) would only involve a power penalty of some 0.7 W (moving the power cost to 2.3 W from 1.6 W) so this exigency might well be worth the additional power consumption if it involves medical or security alarm considerations. The alternative would be to lever the small MCU memory footprint to store the necessary routing information to facilitate this action. The profile absent energizing the main memory looks as follows:

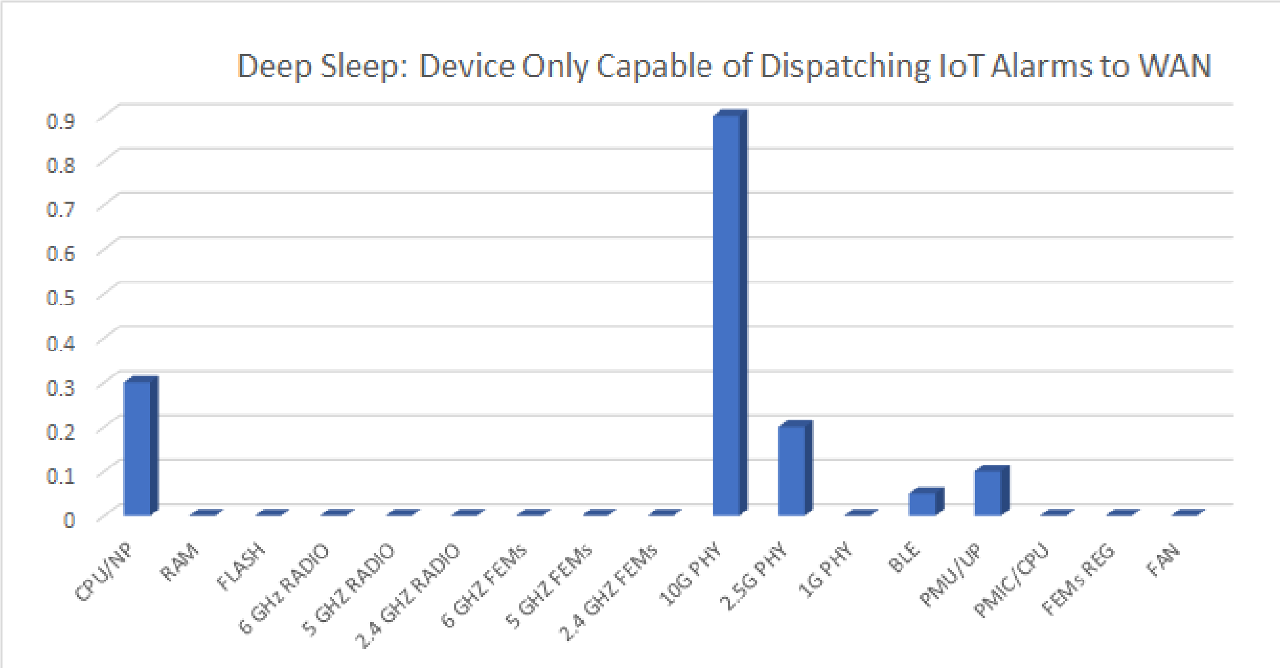

Figure 9 – Deep Sleep Power Profile (Memory Off), 1.6W

Finally, we have the near-fully-lobotomized variant of AP operations associated with the "Off" mode. This fundamentally recreates the ToD-driven brute-force ON/OFF control described in the introductory notes and involves a reconfiguration of the device for it to migrate to, for example, the standby mode. This reconfiguration would involve container image reprofiling so that all ancillary IP services (those sitting astride flows, for example) would be re-instantiated. The exit from this state to something less custodial (or dormant) in fact represents the ToD point where prior history indicates full AP leverage (or "normal" operation) is imminent. The near-empty power landscape for "Off" looks as follows:

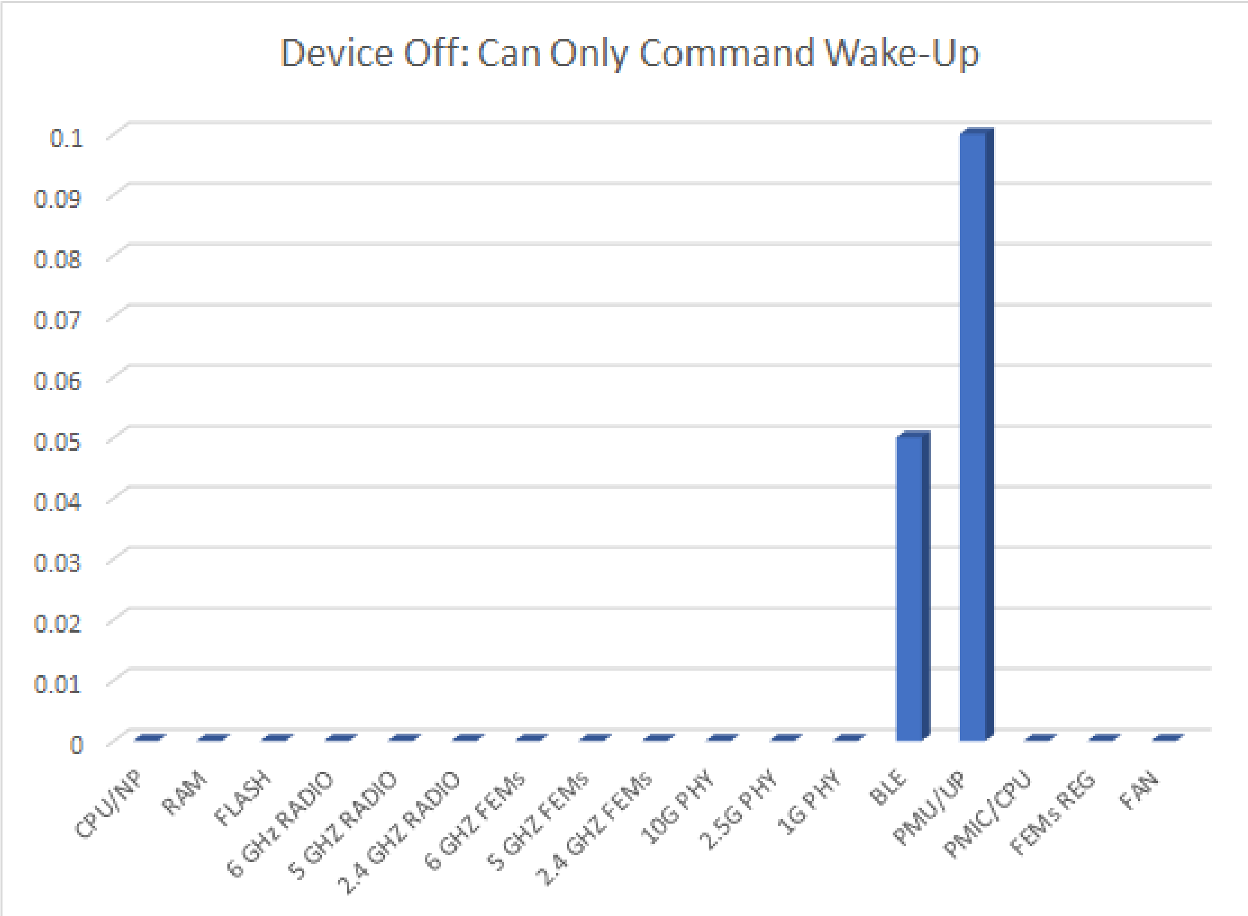

Figure 10 – Device "Off"; Power << ¼ Watt

Note that the IoT is energized such that an application to command power up based upon, say, a cell phone would be possible (to skip past the TOD-driven auto-awakening and allow the user to migrate the device to Standby or Active on a commanded basis).

2.7. Considerations on the Standby -> Active Sweet Spot

Between the 6W of Standby power consumption and the 20W listed for the Active state lie many options for parking the AP's behavior. Below is the power cost matrix which breaks out dissipation of the various circuit functions on a per profile basis:

Circuit Block	Full	Active	Stndby	Dsleep	Off
CPU/NP	2.6	2	1.3	0.3	0
RAM	0.4	0.4	0.4	0	0
FLASH	0.3	0.3	0.3	0	0
6 GHz RADIO	4.3	2.2	1.1	0	0
5 GHZ RADIO	3.3	1.7	0.8	0	0
2.4 GHZ RADIO	2.1	1.1	0.5	0	0
6 GHZ FEMs	5.1	3	0.02	0	0
5 GHZ FEMs	4.8	2.8	0.01	0	0
2.4 GHZ FEMs	4.3	2.5	0.01	0	0
10G PHY	2.8	1.4	0.9	0.9	0
2.5G PHY	1.2	0.6	0.2	0.2	0
1G PHY	0.7	0.5	0	0	0
BLE	0.05	0.05	0.05	0.05	0.05
PMU/UP	0.1	0.1	0.1	0.1	0.1
PMIC/CPU	0.3	0.2	0.1	0	0
FEMs REG	1.4	0.8	0	0	0
FAN	0.8	0.4	0.2	0	0
Power	34.55	20.05	5.99	1.55	0.15

Figure 11 – Power Cost Matrix per AP Functional Circuit

Note that the FEMs represent a frangible line entry, i.e., as the AP has four each spatial streams per band, the power indicated is equally assigned over four devices – each of which may be selectively disabled. This will provide us with additional control granularity for enabling traffic in a particular band, especially as most of the client connectivity represents two spatial streams. (An extender in a mesh network would be a likely four SS link; reduction in spatial density requiring some added consideration since it would impact network performance for several potential clients.) Note that ~7W are tied up in FEM leverage in the Active state, along with 5W in having all bands actively tended to (relative to Standby mode); it is permutation in these areas in particular which will motivate the establishment of certain software behavioral triggers in the following sections.

3. Managing the Power Consumed: Software Behaviors

With the AP's hardware being suitably modified for granular power consumption control, it is now time to discuss how device management can be structured for politic application of assets in various situations to effect responsible standby behavior yet still seamlessly aggregate service demands as these mount. Recall the mention of ToD implications: based on this qualifier, the AP's triggers (and poll delays or time constants) for various asset employment or retirement are bound to diverge as should the expectations for coverage (on average) over the course of 24 hours. Responsivity during peak demand should be necessarily more aggressive at applying a heft of operational assets than might be manifest in device reactionary behavior in periods which experience low average use. Deep packet inspection (DPI) against known service fingerprints also provides the necessary clues (along with historical context) as to the level of QoS (defined along the two axes of bitrate and latency of delivery) expectation the handled traffic has of the LAN.

Note that software administration for the conservation efforts is typically hosted in two domains: AP/local and cloud. The partitioning for these efforts comes by way of offloading situational planning as much as possible into the cloud, with the expectation that this appropriately leverages large-scale machine learning assets in the cloud and increases the level of dormancy in the local device during periods of hibernating service demand. For example, packet differentiated service code point (DSCP) marking for selective QoS handling would already be largely hosted within the cloud, and it then becomes more straightforward to cross this and related packet tagging with ToD and history-related service mount observations and supply the AP with trigger profiles appropriate to when latency-sensitive services are expected to manifest within a given household so that it may be resource-prepped to serve these needs when they arise. (This might involve pre-emptive energizing of the otherwise dormant 6 GHz band when an opportunity arises to mount a highly latency-sensitive service on a client which supports such connection). ToD would also be a bellwether for how deep a hibernating sleep the AP exhibits over the course of the day; maximal retirement of capability would be expected when humans are least likely to seek services of the AP, obviously.

Based on our prior observation of generic ToD usage, a blanket template which sees to lowest power usage from 11 pm to 7 am seems appropriate. The implication here is that, even if one cannot achieve absolute zero power use during this time, such is the goal and one should be at one's most reluctant to apply dissipative resources to any activation during this period; furthermore, gaps in demand should be met with rapid withdrawal of capability since the probability is that any use is likely to be a short, single-link exception to otherwise deep sleep of the AP. (Note that if sporadic history of overnight use is manifest in the device operational history, such may preclude resort of the deepest sleep mode, since this involves the greatest effort at re-initialization of the device upon wake-up – and would perhaps cause more aggravation than would be appreciated by what might be a hard-pressed nocturnal user).

3.1. Band-Specific Characteristics

The three available major bands for Wi-Fi – 2.4, 5 and 6 GHz – manifest different operating environments and serve as residences for differing populations of client devices (and hence, services). 2.4 GHz is the oldest and narrowest band (at 80 MHz total) with the narrowest permissible channels, addressable by all extant MACs and capable of the longest service throws (due to the combination of effective isotropic radiating power (EIRP) and longer relative wavelength) but with the lowest service bitrates – yet also tends towards oversubscription and behaves as a host for overlapping basic service sets (OBSS) and non-WiFi energies (IoT PHYs associated with Bluetooth, Zigbee and Thread). Typically, all Wi-Fi clients have access to this band but because of legacy support, this also implies that a mix of all Wi-Fi MACs can be expected to be encountered (negating many of the significant partial channel and scheduling exploit benefits of the newer MACs – 6 and 7).

5 GHz Wi-Fi also supports all of the MACs 4-7 and features a much wider bandwidth (775 megahertz (MHz), if the seldom exploited U-NII-4 band is included). The shorter wavelength struggles a bit more with signal throw than the 2.4 but there are many 80 MHz wide channels to serve (with appropriate observation and keep-out administered by the dynamic frequency selection (DFS) considerations associated with co-use of some band sections by the military). With a dense enough short/mid-path-length spectral density (64-quadrature amplitude modulation (QAM) and above) and the widest channel BWs, this potentially yields Gbps+ link rates if the stars align on the interference front (meaning no

spectrum puncturing and low contention). Lower latency services like videoconferencing and gaming can park here, along with trunk links associated with in-home Wi-Fi mesh extension. It is more point-to-point than 2.4 GHz but has vastly more capacity in its nearly 10x additional BW and wider channels.

Finally, there is the ultrawide 6 GHz band with its prodigious ~ 1.2 GHz of BW (in the US) and home to multiple 160 and even 320 MHz channels – along with MUCH better channel access determinism due to the fact that access to it is only available to Wi-Fi 6E or 7 link endpoints (and so subject only to their schedule-able MAC uses and not to previous Wi-Fi 4/5 MACs). Even with simple single user (SU), priority-unaware round-robin packet dispatch discipline, links in this band can host low single digit millisecond-grade latencies. The availability of standard power trunk links means multiple-Gbps extension(s) are no problem to instantiate, implying that bulletproof QoS coverage of up to 8,000 square foot (sf) (or even larger) premises for qualifying client populations (those of 6 GHz capability) is easily achieved. In particular, even in contentious environments, the ability to schedule exchanges (absent accommodation for at-will accesses by legacy MACs) mitigates random interferences and means such a time division multiplexing (TDM) mechanism can be paired with the frequency division multiplexing (FDM) flexibility of channel puncturing and fine-grain multi-resource unit (MRU) sub-assignment to greatly improve the usability of very wide channel BWs (160 and 320 MHz) in an extender-based coverage scenario (especially useful in single frequency concatenated meshes). As mentioned before, in early adoption days (until perhaps 2026 or even later) the band will be greatly underutilized and very specifically new client targeted, so less than full exploit – especially under constrained power auspices – ought to be expected. In fact, unless single-band 6 GHz clients specifically emerge (almost all will be dual- if not triband), a transitory scheduling effort at using only the first two of the available 3 Wi-Fi bands (especially in low service demand scenarios) makes good sense from a dissipation mitigation standpoint. (This assumes Wi-Fi 6/7 MAC benefits can be applied to the lower bands where applicable, of course.) This is a bookmark only; The applicability of 6 GHz exploit in a coming algorithmic section will be examined – where service priority will also demand an accounting.

3.2. Wi-Fi Roll of Clients and Services

Recall the previous mention of a budgetary matrix of network capacity/capability versus the roll call of mounted services in the premises and the role this would play in modulating power demand in the AP. Nearly every data-consuming service mounted is hosted by a client, said pairing producing parametric telltales which will define the nature of the link delivery for data to and from the client. The raw poll of parameter assignment precedes an ordering of demand details which will set the band and spatial stream leverages to be applied to the Wi-Fi distribution by the AP – and it is this profiling which will be monitored and instantiated with a bias not just blindly toward maximum performance, but with an eye to conservation of applied power to achieve minimal successful disposition of the ensemble services matrix.

Wi-Fi services have long been categorized under Wi-Fi multimedia (WMM) policies for four priorities of airtime access: audio, video, best effort and background (in descending priority order). We can align more or less with this thinking, bearing in mind that priority will be manifest along the dual axes of bitrate and latency (best QoS associated with highest bitrate and minimal latency). A catalog of a representative mix of service delivery follows:

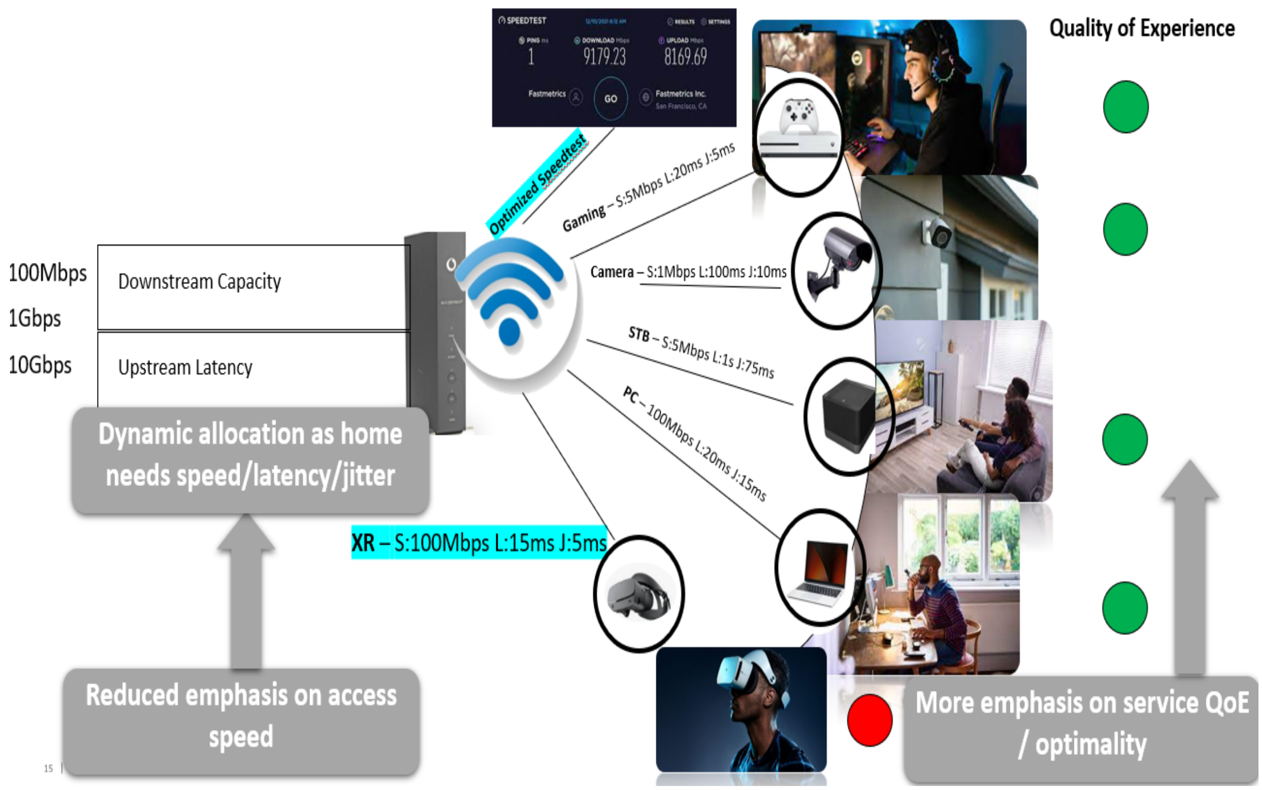

Figure 12 – Mix of Bands, Services and Clients in Wi-Fi Space

In considering delivery requirements, note that both determinism (scheduling conformance to latency bounds) and raw bitrate improve as one advances up in band from 2.4, through 5 and onto 6 GHz. In regard to latency, this can be traced to the reduced occurrence of coincident airtime accesses and an increasing spread of available channels across overlapping basic service set identifications (OBSSIDs) as one moves up-band (allowing channel choice to dodge the worst interference). With respect to bitrate, this evolves in direct reaction to increasing channel bandwidths as one goes up-band. The likelihood of airtime access contention from OBSS or rogue sources also declines over that upwards ramp as well, so ultimately the raw number of services to draw upon the AP will progress along a waterfall where exhaustion of the resource packet delivery "pool" (capacity) available in each band is progressively depleted from a bitrate operating point perspective, thus dictating when the next band up-spectrum is activated to accept some of the load. The pre-emptions to this progression end up being high QoS-based services and client "map-ability" to the various bands, the latter involving exceptions for clients which either cannot connect or connect poorly in some lower band (as in lack of SS resources or, rarely, outright band skip) in deference to operation in a higher band. (In non "green" operations, remember that QoS considerations would dominate the rationale for mapping of band-to-client, with best achievable margins-to-performance the guiding principle and no operational band dormancy tolerated). These two antipathetic strategies actually motivate the potential resort of what amounts to a selectable Econ/Norm switch in AP operational strategy, to permit user input on AP behavior – similar perhaps to the "Econ" mode in some vehicles which maps fuel conservation ahead of outright engine performance. The selection of "Econ" in the AP sense would tend to pack lower bands and subvert full exploit of SS resources where possible – in order to conserve power; "Norm" would move to spread services and pre-emptively push more latency-sensitive traffic to 6 GHz (when client-supported), exploit multi-link operations (MLO) as the situation permits and *always* lever full SS capability.

In any event, the task at hand involves the establishment of service requirement to band capacity ratings. Typical rating of band performance involves use of maximum PHY rates to describe supported bitrates, but these are only achieved at high EIRP and a client service population in immediate proximity to the AP, with all transfers at maximum allowable channel bandwidth and SS leverage; as such, they are unrepresentative and seldom achievable in "normal" use cases. It is therefore better to examine actual airtime costs for connection of a service to a particular client (even though this client may be mobile and subject to some link operating point variability), which at its core involves link EIRP, path losses, modulation and coding scheme (MCS), channel BW and available SS count. Given that some clients can mount more than one wireless service at a time, the airtime budget cost tagging would tabulate as client/service/band/channel/EIRP/MCS/BW/SS. With this history-based 8-tuple for each IP flow, the onset of any client/service pairing (excepting first-time connections) can be pre-emptively assessed for airtime budget impact to existing (or predicted) service mounts based on prior connection and then appropriate band rank ordering can occur for the extant service roster (who gets assigned where, in other words). There is also value in telemetry which monitors delivery performance vis-à-vis packet losses and roundtrip time (RTT) – which provide reliable metrics on contending wireless energies on the channel in question.

Note that, to this point, considerations have been put in place to characterize only SU packet dispatch from the AP (no MRU subgrouping and no MLO binding – the former available in Wi-Fi 6 and 7 and the latter exclusive to Wi-Fi 7). MRU assignment becomes a key adjunct of multi-user (MU) leverage of wideband single channel operations and MLO coupled with MRU can be thought of as a very rare joker card (at least in the coming couple of years) – most invoked for creation of high-capacity multiband trunk links to Wi-Fi 7 capable extenders (operating at peer standard power, perhaps, and capable of MLO/puncturing exploit) and perhaps otherwise reserved for emergency sidecar carriage of data for those few initial clients capable of exercising the MLO option. In a green operational sense, the most likely utilization of these features would fall (by design) to peak use timeframes (the 5-11 pm timeframe) and would otherwise tend to be avoided under conservation auspices unless and until all bands are energized by the AP. The two aspects to tuck away are: 1) for the next couple of years, the number of Wi-Fi clients at MAC epochs 6E and 7 will be relatively low; and 2) scheduling aside, the more advanced aspects of Wi-Fi 7 (very wide BW with MU operations) are not energy-friendly.

3.3. Band Capacity Trigger Points

Each Wi-Fi premises will manifest a history of wireless BW adoption as the active periods progress from early morning to late evening – and such will bear a statistical fingerprint of which a proper AP power scaling algorithm must advantage itself. But as a sort of generic profile (and to provide context where perhaps statistics are not well founded or exhibit wide variance), it may do to project expectations based upon the suggestion provided by Figure 1 in the introduction. Recall that a relatively smooth and monotonic progression (during the diurnal period) of increasing bitrate consumption is forecast by that chart; marrying this to our "increasing bitrate pool" observation above suggests a progression of band leverage that goes as 2.4 GHz only, 5 GHz only, 2.4 + 5 GHz, 5 GHz + 6 GHz (if such step is client-supported) and finally all bands. Note that this progression acknowledges no client band mismatching exceptions and so must be taken with a grain of salt. Accepting this guidance, one can nonetheless calculate band capacity under certain assumptions and generate expansion/collapse trigger points for AP band use. Further stipulation includes that the channel BWs to be used for capacity calculation will be 20 MHz at 2.4 GHz, 80 MHz at 5 GHz and 160 MHz at 6 GHz, with spectral density of no more than 5 bps/Hz (presuming average path losses and median MAC epochs across all subscribed clients) and an airtime transmit density maximum of 25% for 2.4 GHz, 50% for 5 GHz and 90% for 6 GHz (to maintain some margin against transmit backoff accumulation and related deferrals/collisions/OBSS energy – and account for the improvement in this contention behavior as one moves up-band). SS setting at 2 seems advisable as well (though this may be a slight overreach in cases where 3 or less SS are employed by the AP in linking with 2SS-equipped clients). Measurements in the Commscope Wi-Fi house, for example, suggest when spatial streams match, the effective goodput rate essentially reaches 70% of maximum; this translates to 95% or better if a 4x4 is linking with a 2x2 device. The "realistic" band capacities in the generic case then end up being the following:

Band	BW	Capacity
2.4 GHz	20	50
5 GHz	80	400
6 GHz	160	1440

Figure 13 – "Realistic" Channel BW (MHz) and Bitrate Capacities (Mbps) for 2SS Clients from a 4SS AP

Note that, as SWAGs go, these are extremely tenuous dart throws and historical data establishes by far a better – and more premises-appropriate -- view of the wireless landscape and the path losses associated with actual client location distribution. It is not uncommon – in some multiple dwelling unit (MDU) instances especially – to discover that the raw level of OBSS contention at 2.4 GHz renders the band nearly useless for all but background and best effort traffic. It is also true that old "unicorn" single stream clients located far from the AP in any premise can wreck the basic service set (BSS) airtime aperture for other clients in the 2.4 GHz band. As regards to the airtime density factor, this is reasonably inflated (so: improved airtime accesses) as one goes up-band due to MAC maturation and improved channel location options in those instances. Nonetheless, these estimates give us rough per-band bitrate pool capacities we can consider for band energizing triggers as AP utilization ramps up over the course of a day.

Given this stepped view of bitrate demand as posed by the daily cycling, we can perhaps divide the daily cycle into 4 "bins":

Period	Start	End	Duration
Start Up	700	1200	5
Mid-day	1200	1700	5
Evening Pk	1700	2300	6
Night/Idle	2300	700	8

Figure 14 – 24-hr Partitioning of AP Operations (hours)

Referring to our band progression mentioned above, this would seem to drive a band-to-ToD mapping (at least in regard to the Active portion of the AP's operations, between 0700 and 2300 hours) if one, at first cut, regards all clients as at least dual-band capable (and then deals with exclusions as they arise). In addition to the progressive band energizing, there will also be circumspection in assigning full SS resources to a given band until nearing the pool limit (or being driven there by packet loss telemetry which suggests a move to marginally higher capacity is warranted). A boilerplate migration (with SS options) most appropriate to early 6 GHz days (so minimal posed unique-to-band traffic) would look as follows:

	Time	Time	Hrs		Bitrate	Power	Bitrate	Power
Period	Start	End	Duration	Band	w/2SS	w/2SS	w/4SS	w/4SS
Start Up	700	1200	5	2.4	35	10	50	13.2
Mid-day	1200	1700	5	5	280	10.9	400	14.9
Evening Pk	1700	2300	6	2.4+5	315	14.2	450	20.7
Night/Idle	2300	700	8	none	0	0.2	0	0.2
					Mbps	W	Mbps	W

Figure 15 – Boilerplate Bitrate vs Power for AP Operational Permutations over 24 hrs

Note that, for purposes of initial analysis, leverage of the 6 GHz band is not done at all here. It is also presumed that ToD (or historical data) will provide an opportunity to more deeply hibernate the AP, and then awaken it prior to the start-up period (energizing the main CPU/system-on-chip (SoC) and memory and taking the time to establish the correct container profiles). The suggested transitions above provide some permutations as regards SS leverage and so those options (with bitrate benefits and power cost) are shown. In the light-duty case, where all traffic is resolvable to 2SS service only on progressive employ of the 2.4 and 5 GHz bands, one ends up achieving an average daily power use of 191 Whr. For perspective, an unmanaged AP over the course of 24 hours would consume just over double that amount, at around 390 Whr (assuming it can collapse to no better than 6W standby during OFF time). An interesting telltale could be provided by AP's self-reporting "Norm" and "Econ" mode power data over the course of 24 hours, providing users with some value guidance (if they wish to perform the experiment for their particular premises).

3.4. But What About 6 GHz?

The previous section detailed all the reasons for eschewing leverage of the 6 GHz band until some critical mass point of 6 GHz client service is (in future) involved; it turns out that this trigger point happens quickly and for very good reason. If one takes our prior figure on "realistic" bitrate capacity on the various bands (to see when one might begin energizing more bands based on expected capacity exhaustion over the course of a day) and add in the power cost for powering up the band exclusively with 4 SS worth of spatial coverage, a calculation that amounts to a transducer gain (as Mbps/W) worth of band efficiency and the result provides a startling modification of our boilerplate plan can be made:

Band	BW	Capacity	Power	Bitrate/W
2.4 GHz	20	50	10.35	4.83
5 GHz	80	400	11.25	35.56
6 GHz	160	1440	11.95	120.50
			W, 4SS @	Mbps/W
			50%	

Figure 16 – Relative "Realistic" Band Efficiencies, Bitrate/W

The blunt takeaway is that catering to 2.4 GHz clients on behalf of their legacy installation costs dearly on the power efficiency front and instead operation of that band ought to be considered the least desirable option (a huge opportunity power cost) for any service. This observation underscores the true cost of continuing to use 2.4 GHz-only clients in a modern Wi-Fi mesh; add to this the difficulty of providing decent low latency service in the band due to its multi-PHY oversubscription and narrow spectrum/channels and perhaps the best one can do for users in pursuit of AP energy efficiency is to have them substitute newer 5- or 6-GHz clients for older 2.4-GHz devices.

So, what does this do to the boilerplate approach? Essentially, it puts three bookmarks in place right away: if there are no exclusively 2.4 GHz clients in the home LAN, one should move initially to energizing only the 5 GHz band for Wi-Fi use instead of the 2.4 and move the dual- and tri-band clients to 5 GHz service. Second, one should consider 2.4 GHz only for exclusive-to-band clients and those whose service radius otherwise exhausts the AP's throw. (Remember that, on the services front, backhaul of 2.4 GHz-based security camera feeds likely hits both of these parameters). Third, to the extent that both lower bands *must* be energized anyway, it becomes worthwhile to query how much of the day's capacity demands can be met with only 2SS on each of the first two bands (recall the "realistic" estimate that this would serve up to 315 Mbps or so worth of bitrate – though, of course, admittedly perhaps not to acceptable latency limits).

All of this gets to the question of 6 GHz, succinctly put: "when should it be used in a power conserving paradigm?" The answer becomes a calculation on opportunity cost; given that new FEMs almost exclusively only dissipate in transmit mode (very much reduced standby operation power), the 6 GHz band should be energized when the service load power cost undercuts operational power cost in an alternate band by an amount greater than the cost associated with 6 GHz radio chip dissipation – approximately 2 W. Based on our previous calculations for band-specific bitrate-to-power efficiencies, this occurs at approximately 101 Mbps for a 6 GHz-based replacement of a 5 GHz-based service and only 10 Mbps for a 2.4 GHz-based service (though it would be rare to find a client with only 2.4 and 6 GHz connectivity – and those multiband capable should have already been migrated to the more efficient 5 GHz band). The algebra looks as below:

$$\frac{\text{Bitrate}_{\text{SERVICE}}}{\text{Transducer}_{x}} >= \frac{\text{Bitrate}_{\text{SERVICE}}}{\text{Transducer}_{6}} + 2$$

$$\vdots$$

$$\text{Bitrate}_{\text{SERVICE}} >= \frac{2 * \text{Transducer}_{x} * \text{Transducer}_{6}}{\text{Transducer}_{6} - \text{Transducer}_{x}}$$

$\text{Bitrate}_{\text{SERVICE}}$ ➡ Necessary amount of service bitrate to motivate the 6 GHz band use

Transducer_{x} ➡ W/Mbps for band x

Transducer_{6} ➡ W/Mbps for 6 GHz band

Figure 17 – Tipping Point Calculation for Xfer of Service from Band X to 6 GHz

(Quick sidebar: the same tipping point calculation can be done for 2.4 GHz to 5 GHz loading using a 1.5 W offset factor instead of the 2 W above and the result ends up being 8.4 Mbps; this underscores the comments above that the 2.4 GHz band is best used for 2.4 GHz-only client legacy work – and after that, really suited only for best effort and background LAN tasks for DBC/TBC clients as a sort of emergency offload. As soon as service bitrates approaching this low Mbps are reached, it makes more sense from a power-efficient packet delivery standpoint to immediately move the service to a higher band if the client supports it.)

So now there are some directions for the considered exploit of the 6 GHz band in a power-conserving paradigm. First, if one has 6-GHz band only clients, there is obviously no choice but to enable the AP’s use of the band. Second, once one has exhausted the available 4SS capacity pool of 2.4/5 GHz for multiband clients with 6 GHz capability, there is need to transfer some of this load (overflow) to 6 GHz to those clients capable of leveraging it. Finally, even if the pool capacity has not been exhausted for 2.4/5 GHz services but such have accumulated either 101 Mbps worth of 6 GHz client service offload of the 5 GHz band – or (unusually) 10 Mbps of service offload from the 2.4 Ghz band (for, say, a triband-capable client), then the 6 GHz band should be enabled, and these (and subsequent band-matching) service prospects moved to it. Note that this offload applies equally for trunk link replacement – so while 100 Mbps seems like a fair amount, the onset for a 5 GHz trunk link replacement (assuming the extender in question is 6 GHz-capable) can happen quickly.

One final note on 6 GHz exploit and this has to do with considerations on an alternate axis from raw bitrate support: latency. Unlike the determination of opportunity cost on the bitrate capacity front, the consideration for latency typically comes down to two less calculable considerations. The first is an absolute LAN latency limit requirement of less than 5 msec or so; if this is the case (as would be true for perhaps gaming or related simulation work) and the service has a 6 GHz mount option for its leveraging client, then such should be chosen. The second consideration is based on telemetry alarms regarding QoS SLA in a dual-band concurrent (DBC)/TBC scenario, where the SLA is being routinely violated despite even 4SS connectivity to the client in any other band option; in this case, power considerations must be sidelined and if the client is 6 GHz-capable, the band necessarily must be enabled, and the affected client/service pairing moved to 6 GHz to (hopefully) resolve the QoS violation.

3.5. Proposed Flow for Implementation of AP Power Management

It is now time to describe behavior for power managing a TBC AP; the hardware is suitably granular in control and candidate leverage points exist along with an understanding of the merits and parametric weights associated with activation of these controls – enough at least to propose a general flow diagram which addresses shifting cyclic service loads over the course of their (daily) frequency of occurrence. As previously noted, the guidance here will not be to maximize performance unilaterally but to attempt to scale the AP's power use to acceptable levels of accommodation for the services which progressively draw higher bitrates from early morning until late evening of every day. The principal drivers for AP change will be ToD, client population band dependencies, service roster, QoS SLA(s), premise service history, efficiency of the various bands and link telemetry of ongoing transfers. In general, ToD would be used to start and stop the AP in alignment with historical patterns for start/stop (for the latter, the presumption would be that the traffic would have already stopped; otherwise, the ToD STOP becomes an arming signal for the Sleep process which awaits a timeout period of inactivity before commanding shutoff). The proposed control block diagram goes as follows:

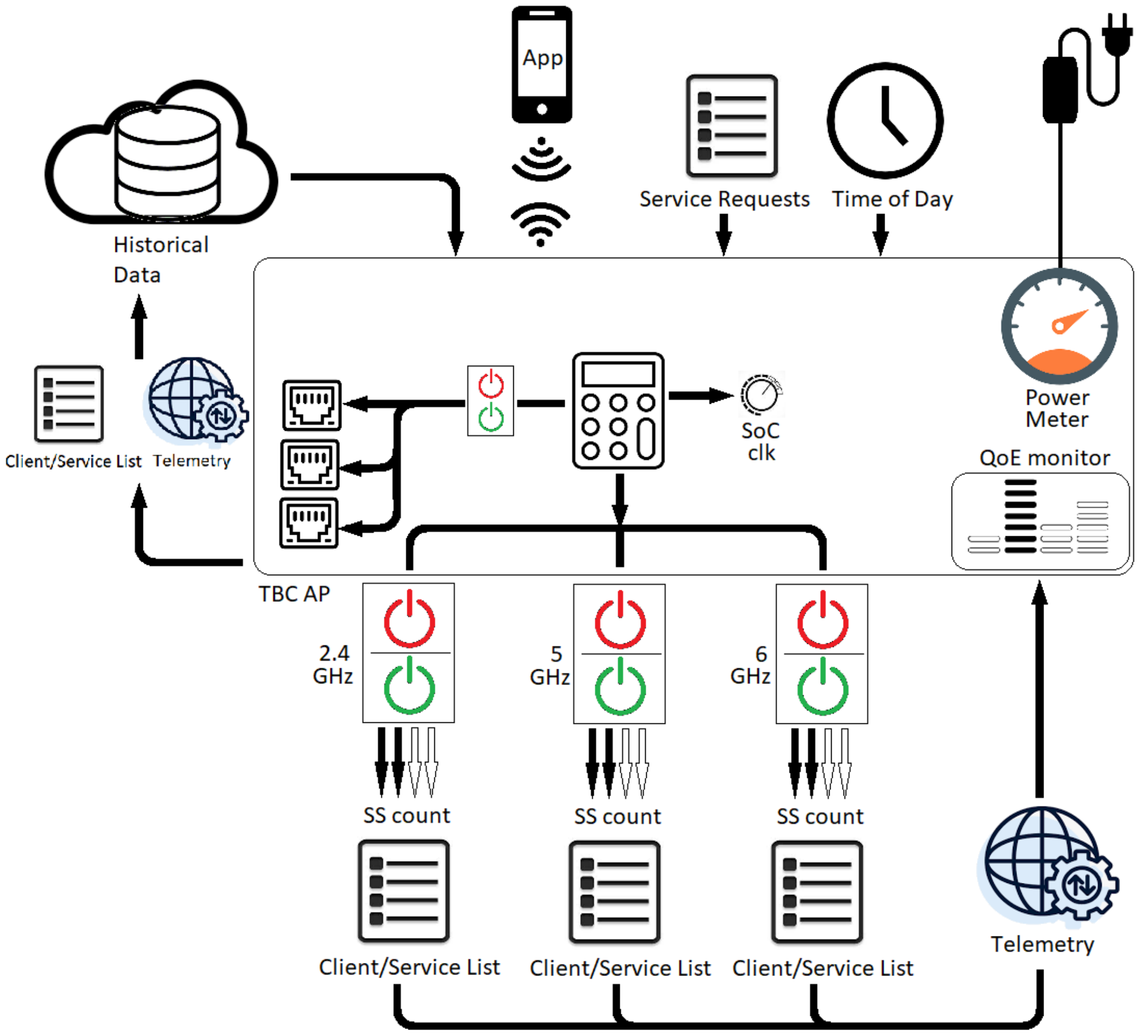

Figure 18 – AP Block Diagram Indicating Green Operations

The general flow of decision-making for transit of the above block would consist of time-driven ON and OFF cycles and then operations within the ON domain. The general ON/OFF block flow (which adjusts AP state based on ToD and history) goes as follows:

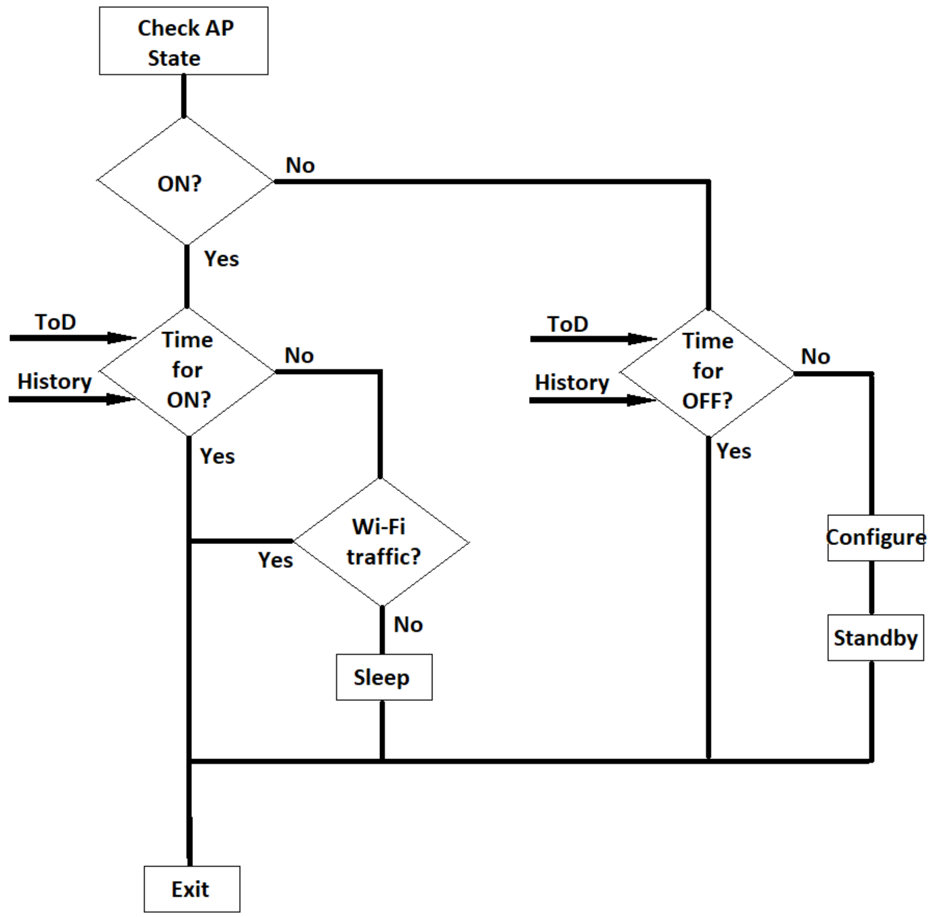

Figure 19 – Polling Logic to Turn AP ON/OFF

Note that, as referred to previously, a phone application over BLE could be used to interdict the automated pacing of AP availability and selected profile or mode, either to pre-emptively awaken the device ahead of standard schedule, change the Econ/Norm mode or to extend operating hours (as either a day extension or an abrupt, atomic nocturnal event). The value of storing history over a week-long period (or longer) allows the AP to adjust to weekend schedules, where daily work usage patterns (and the related service profiles) might find themselves significantly altered from midweek norms.

The next illustration addresses the regular diurnal progression from Standby through end-of-day operations; the summarized operation strategy looks as follows:

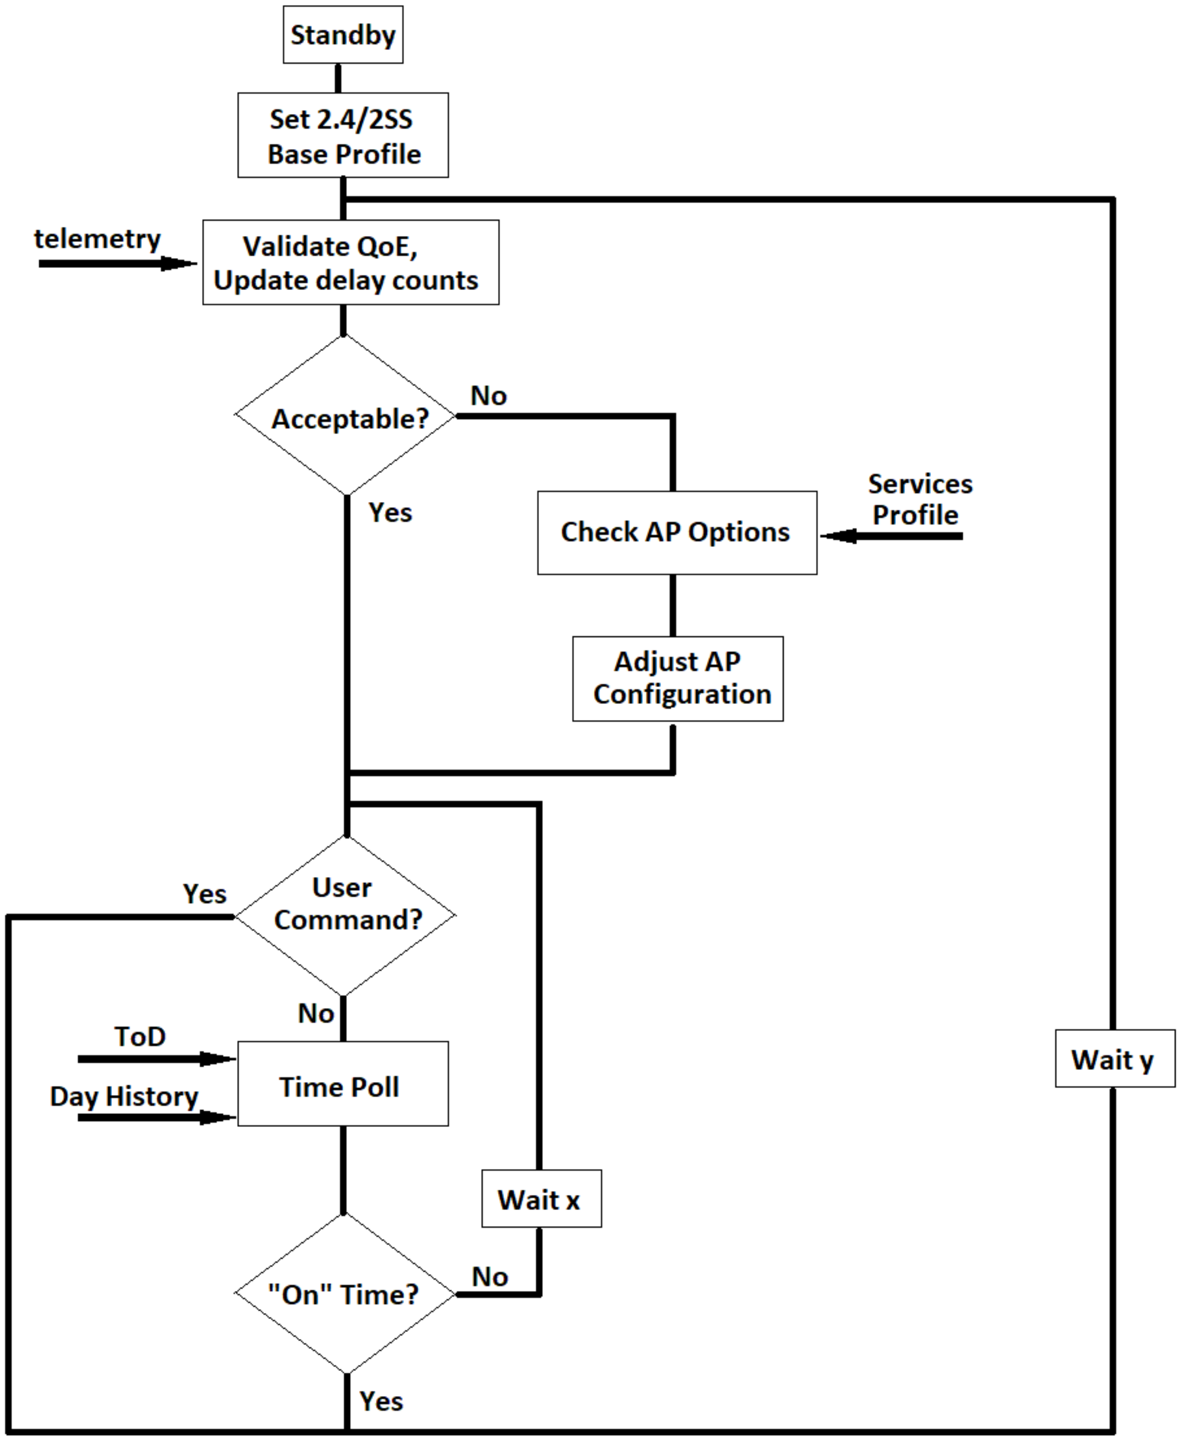

Figure 20 – Moderation of AP Profile to Fit Services Mounted

The hold times for the various AP states can work against "speed to upgrade" versus "speed to downgrade" considerations associated with ON or OFF state and the ToD. Once committed to 4SS service on a given band during diurnal operations, for example, the AP could be set to only slowly decommission some of the spatial coverage, pending expected resumption or expansion of day services. In inverse consideration, if 2SS is already committed to a given band and hints of QoE miss are noted, upgrade to 4SS quickly may be advisable (as would energizing, and moving to, an available up-band to at least 2SS, if such matches the client band coverages).

The progression of 2.4, 5, 2.4+ 5, 5+6 and all bands – subject to client band matching and tipping point calculations (band efficiency fitness for services large enough to be offloaded from lower bands) – provides the basic guidance in these determinations.

4. Conclusion

The European Code of Conduct regarding expectations for CPE provides guidance on what expectations for responsible powering of APs should target; and while US cost of electricity undercuts the European markets by up to 3x (tending to mute the energy responsibility outlook domestically), the opportunity for homogeneous design elements (including SoC functional partitioning) in product sold on a common basis around the world, coupled with shared appreciation for similar specification, seem likely to eventually land the majority portion of these eCOC KPI into US Energy Star profiles. Anticipating such, pre-emptive SP endorsement and patronage of industry efforts to reduce the energy profile of deployed CPE network endpoints would almost certainly experience market reward for responsible citizenship. In supporting task-appropriate energy consumption, the AP signals that it belongs in a 21st-century roster of responsible home electrical appliances – perhaps a bit late to the electrical use sustainability party, but ultimately a conscientious and clever device in that regard.

Abbreviations

3D	three dimensional
5G	fifth generation
AC	alternating current
ADC	analog-to-digital converter
AP	access point
BLE	Bluetooth low energy
bps	bits per second
BSS	basic service set
BW	bandwidth
CPE	consumer premises equipment
CPU	central processing unit
DBC	dual-band concurrent
DC	direct current
DPI	deep packet inspection
DSCP	differentiated service code point
eCOC	European code of conduct
DFS	dynamic frequency selection
DMIPs	Dhrystone millions of instructions per second
EIRP	effective isotropic radiated power
FDM	frequency-division multiplexing
FEMs	frontend modules
Gbps	gigabits per second
GHz	gigahertz
HW	hardware
Hz	hertz
ID	industrial design
iFEM	internal frontend module
IoT	internet of things
IP	internet protocol
KPI	key performance indicator
LAN	local area network
LoRa	long range, wide area
LTE	long term evolution
MAC	medium access control
Mbps	megabits per second
MCU	microcontroller unit
MCS	modulation coding scheme
MDU	multi-dwelling unit
MHz	megahertz
MLO	multi-link operation
MRU	multi-resource unit
MU	multi-user

NFC	near field communications
OBSS	overlapping basic service sets
OBSSID	overlapping basic service set identification
PAN	personal area network
PHY	physical layer
QAM	quadrature amplitude modulation
QoE	quality of experience
QoS	quality of service
RAM	random access memory
RF	radio frequency
RTT	round-trip time
SAR	segmentation and re-assembly
sf	square feet
SG	service group
SLA	service level agreement
SoC	system on a chip
SP	service provider
SS	spatial streams
SU	single user
SW	software
SWAG	scientific wild-a$$ guess
TBC	tri-band concurrent
TDM	time-division multiplexing
ToD	time of day
TXOPs	transmit opportunities
W	watt
WAN	wide area network
Whr	watt-hour
WMM	Wi-Fi multimedia

Bibliography & References

Broadband Capacity Growth Models, John Ulm, Zoran Maricevic, Ram Ranganathan, SCTE Technical Paper 2022

