Power Optimization for Sustainable Smart Speakers: Echo Pop Case Study

Riccha Tripati
Hardware Power & Architecture
Amazon Devices & Services
Sunnyvale, CA USA
Email: rtripati@amazon.com

Dibyendu Nandy
Hardware System Architecture
Amazon Devices & Services
Sunnyvale, CA USA
Email: diby@amazon.com

Marco Pulido
Hardware Power & Architecture
Amazon Devices & Services
Sunnyvale, CA USA

Joey Escobar
Hardware Power & Architecture
Amazon Devices & Services
Sunnyvale, CA USA

Jeongwoo Han
Sustainability Science
Amazon Web Services
Seattle, WA USA

***Abstract*— In this paper, we present a comprehensive system-level approach to advancing device sustainability through power optimization for smart home devices, with a detailed case study of Amazon's Echo Pop. Through Lifecycle Assessment (LCA), we identified that Echo Pop generates an estimated 42 kg CO2e over its product lifetime, with 24 kg CO2e (57%) attributed to use-phase emissions, highlighting the critical importance of idle power optimization for decarbonization efforts. We implemented a novel system architecture for energy efficiency that leverages CPU Suspend-to-RAM states and Wi-Fi power save modes. We minimize energy consumption during device inactivity and via coordinating various system services to maintain seamless user experience. The system intelligently transitions between power states using a low power Digital Signal Processing (DSP) core for monitoring ambient audio, while duty-cycling background connectivity tasks.**

We validated the effectiveness of our multi-domain power management approach across SoC, Wi-Fi, and system-level components, through comprehensive power rail analysis using high-resolution measurement methodologies in a lab setting. Results demonstrate substantial power reduction achievements, with Echo Pop achieving approximately 1.1W standby power consumption—a 49% improvement over previous generation devices that drew over 1.5W. Ground truth validation through in-field telemetry data confirms our lab projection models, with a daily average energy consumption of 28.21Wh closely matching the in-field empirical measurements of 28.07 Wh. This work establishes a validated framework for sustainable smart device design that balances environmental impact reduction with maintained functionality and user experience.

Keywords*— *low power

I. INTRODUCTION

This paper presents Amazon Devices' comprehensive approach to energy-efficient design as part of its broader commitment to the Climate Pledge, sustainable shopping program, and corporate Net-Zero carbon goals [1].

The energy consumption of smart home plugged-in devices during their operational lifetime represents a significant portion of their total carbon footprint, as revealed through Lifecycle Assessment (LCA) methodology, a scientific framework used to evaluate energy use, emissions and wastes generated from direct and indirect processes. LCA is a widely accepted method to quantify environmental impact of a product starting from Materials and Manufacturing, through life time Use till End of Life; including climate impact measured in kg CO2 equivalent (kg CO2e) [2-4]. Devices' idle power consumption is of particular importance for reducing use-phase emissions.

Numerous techniques and strategies have been proposed to reduce smart devices' energy consumption while maintaining the quality of services. Uddin and Nadeem (2013) proposed A2PSM [5], an audio-assisted Wi-Fi power saving mechanism that uses inaudible high-frequency audio signals to reduce Wi-Fi interface wake-ups, achieving 25% greater power savings than standard schemes. Wang et al. (2015) proposed LAD (Light-weight Adaptive Duty-cycling) [6], a distributed protocol for wireless sensor networks that dynamically adjusts duty-cycling parameters based on traffic conditions, achieving 28.2%-40.1% energy consumption reduction compared to state-of-the-art protocols while maintaining high reliability. Castro et al. (2024) demonstrated that dynamic duty cycling in coherent DSP ASICs [7] could reduce power consumption by 22-74% in carrier frequency offset estimator algorithms by adaptively adjusting sleep periods based on real-time operating conditions rather than worst-case scenarios. Recent researches focus on utilizing machine learning (ML) algorithm to enhance devices' energy efficiency further, such as Masood et al. (2024), Wang et al. (2025), Wright et al. (2025) [8-10]. And Rashid et al. (2025) implementing energy efficiency ML algorithm on edge devices [11].

Prior work have demonstrated the effectiveness of the proposed solutions for specific components or tasks. As smart home devices often utilize multiple energy efficiency improvement solutions and infield duty cycles vary widely by users, analyzing the impacts of such solutions in real infield devices is critical to improve devices' energy efficiency further.

This study aims to fill the gap by presenting a detailed analysis of hardware power optimization strategies on Echo Pop, a popular high volume consumer electronics product from Amazon Devices. The Echo Pop has an estimated carbon impact of 42 kg CO2e over product lifetime per unit, of which 24 kg CO2e is contributed by use-phase energy consumption [12]. The LCA assumes that the average product lifetime is 5 years for these devices while actual lifetime of devices varies by users. The 5-year-lifetime assumption is consistent with the lifetime of TVs assumed in EPA's Energy Star program. Through the implementation of lower power states, idle mode power consumption of device has been substantially improved over previous generation Echo products and is competitive with similar consumer electronics in the industry.

This paper explores system level energy reduction, measurement methodologies for sustainability focused integration. We examine the challenges and opportunities in energy characterization and optimization while considering upstream dependencies across silicon design, software, and system architecture layers. The energy consumption projections are validated against empirical metrics reported by devices' ground truth telemetry.

II. Device System Architecture

Energy usage of wall-plugged devices are a carbon hotspot, about 60%, in the footprint of Amazon Devices and offers the biggest single opportunity for carbon reduction. We focus on reducing energy usage due to continuous sensing mode operations for functionalities like wake word detection and wireless connectivity maintenance.

This is done by detecting periods when the device is not in use to seamlessly enter a Low Power Mode (LPM) of operation and exit LPM on demand with minimal user perceived latency impact [13]. Echo devices in a quiescent state are functionally always on to monitor the environment with acoustic sensing and through other available sensors. LPM is designed to opportunistically minimize power consumption without impacting user experience, so when an input triggers a wake the device is brought out of LPM. A key trigger is the sensing and detection of ambient sounds that meet a designed "wake" criteria. Other wake triggers include Ambient Light Sensor changes for adjusting LED brightness, and accelerometer trigger inputs for tap detection. A third set of wake triggers are based on connectivity requirements.

The biggest consumption of power for Echo devices during inactivity is by the SoC (system on chip) and Wi-Fi connectivity chipsets. The power requirements during inactivity derive from monitoring ambient audio in the room to recognize the Wake Word (WW) and monitoring the other sensors on the device. Additionally, devices have background activities to maintain cloud connectivity, such as cloud directives for smart home and speaker features including music playback, Whole Home Audio, peer to peer requests for Drop-In VoIP calls, Notifications, Reminders, Announcements, BT music playback, and similar demands. Devices are required to also run background activities such as time synchronization, maintaining cloud HTTP connection, connectivity for smart-home IoT device control.

The architecture is designed to minimize power consumption when the devices is not engaged with user like an Alexa query, smart home change, or music playback. The architecture minimizes active residency for each sub-system during background activities to the minimum required sub-set and resumes full LPM when wake triggers are absent for a pre-determined hysteresis period and no background activities are being run.

Echo device architecture uses a CPU core Suspend-to-RAM (STR) state for LPM implementation within SoC. Device software is optimized to enable CPU STR for LPM such that required CPU services could be aligned in time to operate when the CPU are periodically brought out of suspend, using timers and allowing service processes to hold wake-locks. Energy saving benefits are most significant by designing power saving modes for both SoC and the Wi-Fi throughout periods of inactivity; lower power states can be entered with minimized connectivity and background activity by running subsystems in a duty cycled manner. Device can exit LPM with minimal latency, and be able to resume any active use case when a wake lock is generated. Performance of background activities, persist across inactive periods.

III. State Machine Definition

Simplified system power states are defined here from the user context for clarity, as Active and LPM. Additionally the Active state is broken down into Active.Standby for background housekeeping activities transparent to the user and Active.Use to indicate that the device operates in a user engagement use case. LPM is a new state in this description that did not exist in prior Echo devices. Figure 1 shows this power state diagram.

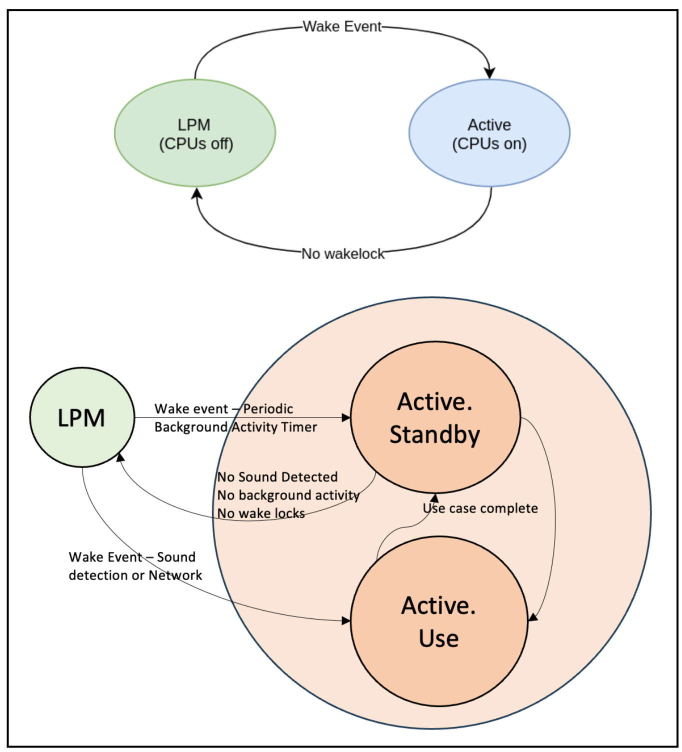

Fig. 1. (a) Power State transition for CPU from Active to Low Power Mode with wake triggers and (b) detailed Power State Machine diagram.

In LPM, power is minimized by leveraging a Digital Signal Processing (DSP) core for monitoring audio and stopping all background tasks on the CPU, so it can enter STR state. Background tasks are duty-cycled to those limited periods when the system periodically enters Active.Standby mode or resumes

Active.Use modes. Active.Use mode refers to device operation required for user interaction, such as WW detection for an Alexa query, music playback or other feedback to the user, presence detection through sounds or ultrasonic triggers, or events such as Drop In, Notifications, Alarms, Announcements, and the like. This user oriented states are typically for use cases which require significantly more power than the standby activities which do not directly define user interaction.

A. SoC Power Modes

As stated above, power can be minimized when CPU cores enter STR. LPM must operate transparently without the user being aware when device moves into or out of this mode. For this a fast resume of the full SoC is necessary. CPU cores are power gated. DSP is active to handle audio input frames from the Mic. If it detects a WW or sound packet that indicates user interaction, it will generate wake event trigger for CPU.

The SoC has the following subsystems shown in the system block diagram below, which can be controlled by appropriate voltage domains and levels to meet LPM, Standby and Active requirements. The SoC power domains are based on Multi-Threshold CMOS (MTCMOS) structured under a dedicated power rail for CPU, and separate one for the remaining subsystems including DSP. The CPU domain has its own power rail and also offers fine-grained control and Dynamic Frequency and Voltage Scaling (DVFS) operating modes to achieve high performance and low power in the CPU subsystem, while the VCORE domains are coarse grained and are power gated by MTCMOS switches. The power tree, also shown below, illustrates powered and power gated sub-systems.

B. WiFi Power Save

Connectivity power is a significant part of the system's power consumption, so the Wi-Fi chipset is put into power save mode (PSM) when there is no data traffic. It wakes at periodic intervals for DTIM (Delivery traffic indication message) exchange with the Wi-Fi Access Point.

C. Exit conditions

When device is inactive, it will enter LPM state, but it is designed to exit LPM expediently and with very low latency for user initiated activity to maintain the best possible user experience.

As an example, CPU cores are brought out of suspend by an interrupt from the DSP upon detection of audio indicative of human speech or presence. The same interrupt will also bring the Wi-Fi out of PSM, so that active cloud connectivity is established before WW detection and on-device processing is complete. The device maintains a similar protocol for cloud directives to support cloud directives such as notifications, drop-ins and other user facing features. The device also exits LPM periodically for background process activities, such as maintaining a designed duty cycle for Wi-Fi power save mode using DTIM and the CPU to allow time synchronization for whole home audio or communication of device states for Smart Home Control.

Additional impact is not apparent to the user as these entry and exit operations are implemented with minimal latency by optimizing Wi-Fi DTIM period and CPU resume time.

IV. Power Evaluation and Analysis

Echo devices are powered by AC wall plug with an efficient Power Supply Unit rated for DC 12V. Device lower power can be observed during quiescence, by simply plugging in AC wall adapter to an off the shelf power meter (HOBO plug load logger) and left to idle. Data captured in this way at a sampling rate of 1 sample/second is shown in figure 2.

Further bench tests were conducted on Echo development devices, by breaking out DC system voltage rails, to deep dive on SoC and Wi-Fi behavior and validate design. A system block diagram that graphically represents the on-device sub-systems and their relationships in a process is shown in figure 3. This block diagram visualizes Active.Standby mode and also identifies hardware blocks to optimize for lower power. Within figure 3, a power tree illustrates power delivery network and shows how power flows from main supply, through converters, into specific voltages and currents needed for various loads.

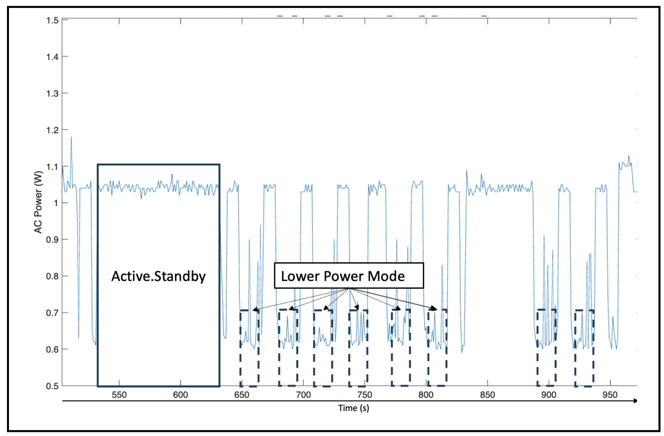

Fig. 2. AC power of device in quiescent state measured with power meter.

A. Measurements

As previously noted the subsystems of interest to monitor for power savings design are CPU, Wi-Fi, and DSP. So the power rails measured to validate LPM are CPU, Wi-Fi, and VCORE. Sense resistors of 0.02 Ohms value were installed on specified rails of Device Under Test to monitor current and differential voltage during using an Analog to Digital Converters (ADC) setup. VSYS power was also captured for total DC system. Data was collected during a typical Active.Standby mode, without any active user engagement, to allow for LPM entry in an acoustically quiet environment devoid of Radio Frequency noise. Power measured per rail at a sampling rate of 2000 samples/seconds with ADC, is plotted in figure 4. It shows a trend of power variation at a high rate due to dependency on device system clocking. Overall, there is power reduction down to 0 on both CPU, and Wi-Fi rails. Whereas the VSYS and VCORE power do not reduce to 0, as the system continues to support the functionality discussed earlier. Further analysis was done to demonstrate the system power dependency on CPU and Wi-Fi power saving schemes.

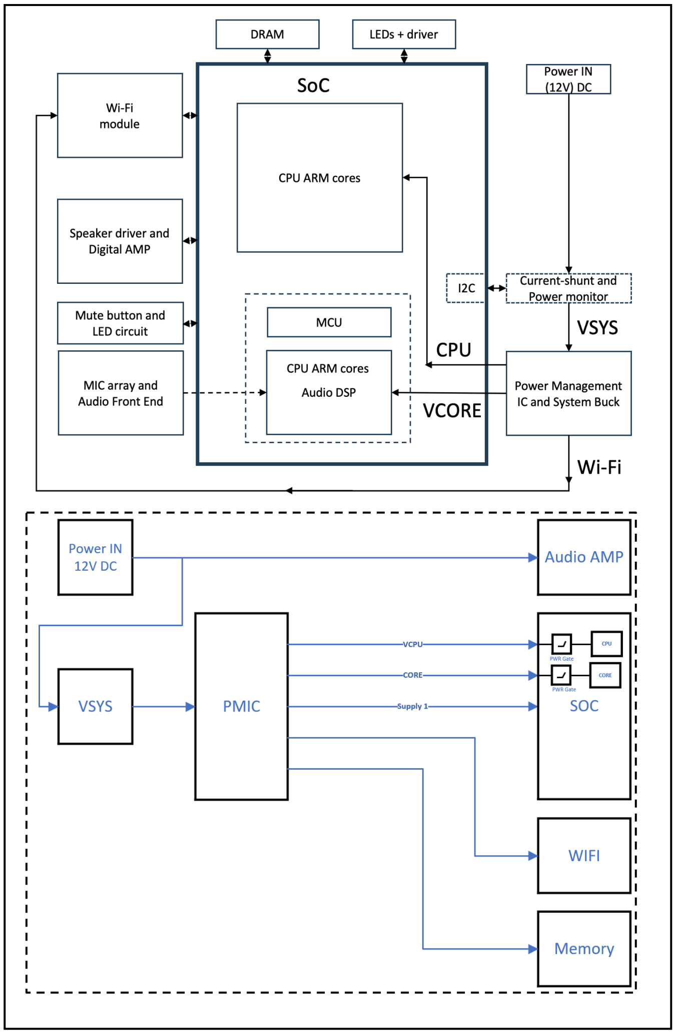

Fig. 3. (a) System block diagram, with telemetry hardware and (b) power tree with power delivery network from power supply input to DC blocks.

B. *Classification*

A classifier function was used to bucket the subsystem rails data from CPU and Wi-Fi into high and low bins. For both power rails the centroid of "low" classified power data is close to 0 validating that CPU clusters do indeed enter STR and Wi-Fi enters PSM, as designed. The classified bins of CPU and Wi-Fi, logic is applied to corresponding VSYS power data point by point at the same time stamps as indicated in table 1.

Time series plot visualizing this is shown in figures 5-6. Notably there were instances of partial LPM observed through data, when only one of CPU or Wi-Fi continued power saving, while the other subsystem remained awake for compute processing or serving connectivity functions. A bi-variate histogram is plotted showing classified power states of CPU and Wi-Fi, to show the time residency distribution per device state in figure 7. Total system power for each state machine was estimated using aggregated samples of binned clusters, by averaging power weighted by time spent in each state

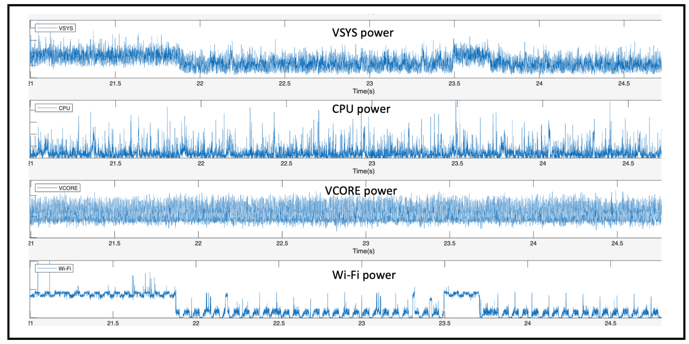

Fig. 4. Time domain system break down: measured DC power for rails: (a) VSYS, (b) CPU, (c) VCORE, and (d) Wi-Fi.

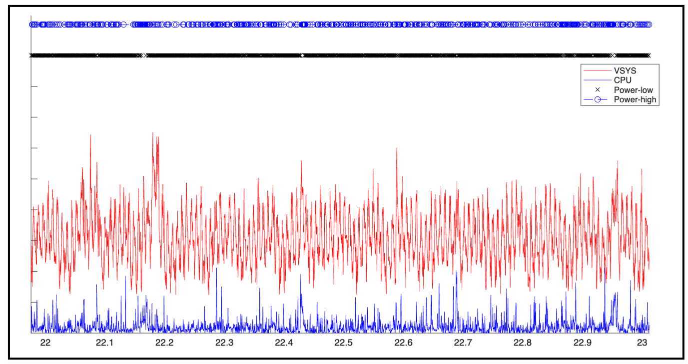

Fig. 5. Time domain classified VSYS points based on CPU power bins.

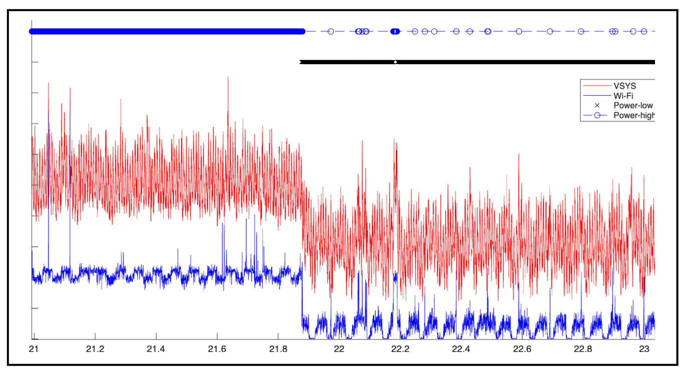

Fig. 6. Time domain classified VSYS points based on Wi-Fi power bins.

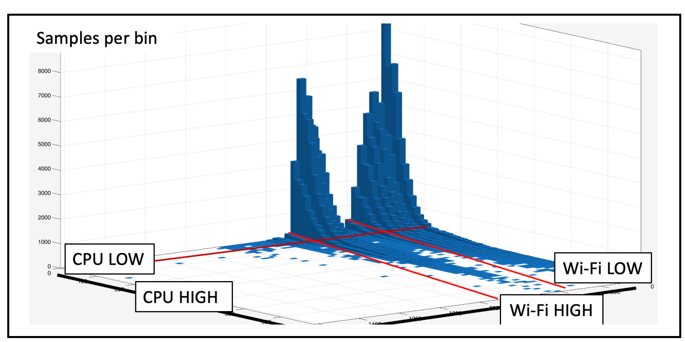

Fig. 7. Bi-variate Histogram of CPU and Wi-Fi power data bins.

TABLE I. LOGIC FOR VSYS DATA CLAFFIER

LOGIC	VSYS power	CPU power	Wi-Fi power
Active.Standby	1.1W	HIGH	HIGH
Partial-LPM	0.7W-0.9W	HIGH xor LOW	
Full-LPM	0.6W	LOW	LOW

C. Ground Truth

These power projections are converted to lifetime device energy consumption by using the formula (1): for plugged-in devices, electricity is consumed at different rates over time as the device alternates between different power modes (e.g., standby, active, LPM). Each mode (*k*) has its own power rating and is measured in kilowatts. The kilowatt-hour consumption for a plugged-in device is equal to the sum of the power rating (P_k) per power mode, multiplied by the lifetime hours in that power mode (T_k), divided by the power adapter efficiency (η)

$$\epsilon^{life} = \frac{\sum_{k=1}^{k} P_k * T_k}{1000*\eta} \quad (1)$$

Some in-field devices are capable of telemetry to capture system DC power daily average using a current sense amplifier, as shown within figure 3. SoC is connected to telemetry hardware through I2C. These devices are reporting a daily average read out by SoC when connectivity is available.

Based on field metrics analysis over one-year, it was found that actual daily energy consumption matched closely with system analysis, as shown in table 2, hence validating that projections represent ground truth. These estimates were converted to use-phase emission equivalent using grid emissions factor [2].

TABLE II. GROUND TRUTH

Energy (Wh)	*System analysis projected*	*Field Metrics reported*
Daily Average	28.21	28.07

V. CONCLUSION

We have demonstrated how Amazon's sustainability initiatives improved energy efficiency of the Echo Pop device from previous generation systems. Echo Pop standby power is around 1.1W or lower due to power optimizations outlined in this paper compared to the previous generation (Echo Dot 3rd gen) that draws over 1.5W in standby. As a result, Echo Pop saves 8.9 kg CO2e of use phase emissions compared to Echo Dot 3rd gen under the same conditions while providing a better user experience. We have been implementing similar techniques across all our Echo smart speaker and Echo Show smart display product lines. In addition to power optimization to drive energy efficiency, Amazon's consumer electronics products utilize recycled materials, recyclable packaging, and less Green-House-Gas intensive transportation modes to lower its overall carbon footprint by an additional 4.6 kg CO2e. These gen-over-gen savings contribute toward Amazon's Climate Pledge. Future directions in our work involve low power mechanisms for efficient presence detection, such HW support for motion detection, optimizing camera frame rates for CV, designing ML architectures for efficient edge AI compute, removing redundancy and optimizing of demanding algorithm workloads and overall software stacks in our devices.

ACKNOWLEDGMENT

We are grateful to Amazon Devices Hardware leadership for supporting innovations presented in this work, and also thank the Devices Sustainability & Science team for shepherding Climate Pledge Friendly certifications of Amazon's devices.

REFERENCES

[1] Amazon, "The Climate Pledge," Amazon.com. [Online]. Available: https://www.aboutamazon.com/planet/climate-pledge. (Accessed: Nov. 13, 2025).

[2] "Amazon Devices Product Carbon Footprint Methodology," sustainability.aboutamazon.com. https://sustainability.aboutamazon.com/devices-carbon-methodology.pdf (Accessed Nov. 13, 2025)

[3] World Resources Institute and World Business Council for Sustainable Development, "GHG Protocol Product Life Cycle Accounting and Reporting Standard," WRI and WBCSD, 2011. [Online]. Available: https://ghgprotocol.org/product-standard. (Accessed: Nov. 13, 2025).

[4] International Standards Organization, "ISO 14067:2018 Greenhouse gases — Carbon footprint of products — Requirements and guidelines for quantification," ISO, 2018. [Online]. Available: https://www.iso.org/standard/71206.html. (Accessed: Nov. 13, 2025).

[5] Mostafa Uddin, Tamer Nadeem, "A2PSM: Audio Assisted Wi-Fi Power Saving Mechanismfor Smart Devices," in *Proceedings of the 14th Workshop on Mobile Computing Systems and Applications*, 2013. doi: 10.1145/2444776.2444782. [Online]. Available: https://dl.acm.org/doi/10.1145/2444776.2444782

[6] J. Wang, Z. Cao, X. Mao, X. -Y. Li and Y. Liu, "Towards Energy Efficient Duty-Cycled Networks: Analysis, Implications and Improvement," in IEEE Transactions on Computers, vol. 65, no. 1, pp. 270-280, 1 Jan. 2016, doi: 10.1109/TC.2015.2417558.

[7] L. Castro et al., "Exploring Dynamic Duty Cycling for Energy Efficiency in Coherent DSP ASIC," in IEEE Embedded Systems Letters, vol. 16, no. 2, pp. 202-205, June 2024, doi: 10.1109/LES.2023.3322301.

[8] F. Masood et al., "AI-Based Wireless Sensor IoT Networks for Energy-Efficient Consumer Electronics Using Stochastic Optimization," in IEEE Transactions on Consumer Electronics, vol. 70, no. 4, pp. 6855-6862, Nov. 2024, doi: 10.1109/TCE.2024.3416035.

[9] J. Wang, Y. Hu and L. Li, "Real-Time Energy Efficiency Optimization in Consumer Electronics via Multi-Scale LSTM and Deep Reinforcement Learning," in IEEE Transactions on Consumer Electronics, vol. 71, no. 3, pp. 7730-7740, Aug. 2025, doi: 10.1109/TCE.2025.3589223.

[10] E. Wright and A. Gopalan, "A Low-Overhead Approach to Dynamic Frequency Scaling for GPUs Using Multi-Dimensional Energy Efficiency Curves," 2025 IEEE Conference on Technologies for Sustainability (SusTech), Los Angeles, CA, USA, 2025, pp. 1-8, doi: 10.1109/SusTech63138.2025.11025714.

[11] H. -A. Rashid, E. Kanjo and T. Mohsenin, "HAC-M-DNN: Hardware Aware Compression of Sustainable Multimodal Deep Neural Networks for Efficient TinyML Deployment," 2025 IEEE Conference on Technologies for Sustainability (SusTech), Los Angeles, CA, USA, 2025, pp. 1-7, doi: 10.1109/SusTech63138.2025.11025634.

[12] Amazon, "Echo Pop environmental sustainability fact sheet," 2023. [Online]. Available: https://sustainability.aboutamazon.com/devices_fact_sheet_echo_pop.pdf. (Accessed: Nov. 13, 2025).

[13] Amazon, "Low Power Mode For Speech Capture Devices," 2022, US Patent US-11514926-B2 & US-11810593-B2, [Online]. Available though USPTO search https://ppubs.uspto.gov/basic