AcouBLE: Clock-Seeding BLE chips via Acoustic Signals for Extended Energy Efficiency

Mohammad Rostami, Gentian Jakllari

AcouBLE: Clock-Seeding BLE chips via Acoustic Signals for Extended Energy Efficiency

Mohammad Rostami and Gentian Jakllari
Univ. Toulouse, Toulouse INP, CNRS, IRIT, Toulouse, France
{mohammad.rostami, gentian.jakllari}@irit.fr

***Abstract*—We present AcouBLE, the first system to significantly reduce the BLE tag's RF front-end energy consumption while preserving standard BLE operation, backward compatibility, and one-to-one deployability. At its core is clock-seeding, a new paradigm that offloads the generation of the high-frequency MHz clock from the energy-constrained tag to the central (smartphone). The smartphone transmits a low-frequency audio signal that the tag transforms into a clean, ultra-low-jitter reference clock–eliminating the most energy-expensive component of a BLE connection event.**

Our design combines a novel digital-analog upconversion pipeline that converts a kHz audio tone into a stable 32 MHz clock with a lightweight feedback protocol to implement clock-seeding on commodity BLE chips without modifying the BLE stack or requiring privileged access on the central. A full prototype demonstrates up to 5.9X energy savings in Connected mode, while maintaining RF performance, seamless deployability, and backward compatibility.

I. Introduction

The Internet of Things (IoT) has rapidly emerged as the backbone of pervasive computing, with over 15 billion connected devices in 2024 and projections reaching 40 billion by 2030 [1]. Bluetooth Low Energy (BLE) is among the most widely adopted wireless technologies in IoT and user-facing devices, accounting for roughly 27% of global IoT connections [2]. BLE's success stems from three factors: its low-cost, commodity hardware; its mature protocol ecosystem supporting diverse IoT applications [3]; and its ubiquitous presence on smartphones and tablets, enabling seamless peer-to-peer connectivity without extra infrastructure. BLE's asymmetric central-peripheral architecture further enhances its practicality. The central device (e.g., smartphone) manages most protocol complexity, allowing the peripheral (e.g., sensor tag, wearable) to remain in low-power sleep for most of the time, waking only for short connection events—a mechanism known as duty cycling.

Yet in energy-constrained IoT devices, such as medical monitors, wearables, environmental loggers, etc., BLE often dominates energy consumption because these devices lack other power-intensive components like displays or high-end processors. For example, BLE radios alone consume over 82% of total power in commercial earbuds [4]. As IoT scales to billions of devices, this disproportionate cost makes BLE energy use a critical bottleneck for pervasive computing systems.

As a result, a growing body of work seeks to improve BLE energy efficiency across multiple layers of the stack. Prior efforts include optimizing protocol parameters [5] (e.g., connection intervals, advertising rates), reducing packet overhead [6], or leveraging low-power wake-up radios [7], [8] to minimize active listening time. While effective, these techniques typically deliver incremental gains, often bounded by the energy cost of the BLE chip's active RF front-end

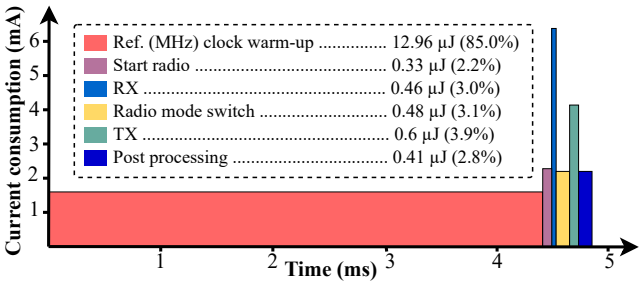

Fig. 1. Energy consumption in peripheral Connected mode of a commodity BLE chip transmitting 10-byte packets. Clock startup dominates per-event energy—85% of total cost—making the clock the real energy bottleneck.

A more radical class of solutions replaces the BLE transmitter entirely with RF backscatter [9], [10], [11], where tags reflect incident signals instead of generating their own. By eliminating active RF transmission, backscatter can reduce communication energy by orders of magnitude, enabling ultra-low-power or even battery-free devices [12], [13]. However, most backscatter systems only support Advertising mode, where tags broadcast short, one-shot packets for device discovery. Only recently [14] has begun exploring backscatter in Connected mode, the standard BLE mode for periodic, bidirectional communication between a central and a peripheral. Yet, backscatter introduces a key barrier to deployability: backscatter systems require specialized infrastructure, such as dedicated exciters or receivers [14], [13], undermining BLE's key advantage of ubiquitous, one-to-one connectivity. As a result, backscatter has seen limited adoption outside specialized research prototypes, despite its unparalleled energy benefits [15], [16].

Problem Statement. Significantly reduce the energy consumption of a BLE tag's RF front-end in Connected mode for periodic, bidirectional communication with a smart device (the central). To be practically deployable at scale, any solution must satisfy two key requirements:

❶ **Commodity-like deployment:** No specialized infrastructure: a one-to-one, central (smart device) to peripheral (IoT tag) deployment.

❷ **Backward compatibility:** Only user-space software (an app) on the smart device and seamless interoperability of the tag with existing BLE devices and legacy systems, enabling incremental adoption.

To tackle this problem, we begin by analyzing the energy consumption of a commodity BLE tag in connected mode across its wake–transmit–sleep cycle (details in § II-A). As shown in Fig. 1, a single connection event, where the tag receives the central's (mandatory) ping and sends a response with 10 bytes of payload, reveals a surprising result: *nearly 85% of the energy is spent not on transmission or reception, but on generating the high-frequency (*MHz*) clock driving the RF front-end*. The reason lies in the startup process of the crystal-based MHz clock in commodity BLE chips: beginning from thermal noise, the crystal oscillator stabilizes through a slow feedback loop, requiring several milliseconds to produce a low-jitter, reliable signal—far longer, and thus far more energy-hungry, than the RF communication itself.

Building on this insight, we propose clock-seeding, a new paradigm that can significantly reduces BLE energy consumption. Unlike backscatter [17], which eliminates the RF front-end entirely, clock-seeding offloads only the high-frequency reference clock generation from the tag to the smart device (central). The smart device, e.g. smartphone, transmits a carefully designed signal that the tag converts into a clean MHz-level clock at sub-μJ energy cost, sharply lowering per-event energy consumption. While it may seem intuitive, clock-seeding raises a fundamental design question: *what signal should the smart device use?* The answer is constrained by the problem statement requirements: ❶ The signal must remain active throughout each BLE connection event, so it cannot share spectrum or hardware with BLE, ruling out options like Wi-Fi. ❷ The signal source must not depend on interfaces that restrict third-party control, such as LTE.

To realize BLE clock-seeding while satisfying requirements ❶ and ❷, we introduce AcouBLE, a system that enables acoustic-aided clock seeding for BLE tags using a commodity smartphone, the natural hub for pervasive computing. AcouBLE makes a surprising but effective choice: it delivers the clock-seeding signal over the acoustic channel. Modern smartphones come with high-quality speakers and user-level audio APIs that allow precise control over sound generation [18]. By transmitting an inaudible tone in the 18 kHz to 20 kHz band, AcouBLE provides a clock-seeding signal that coexists seamlessly with BLE, requires no privileged system access, and demands no hardware modifications on the smartphone.

However, building AcouBLE introduces two key challenges mainly stemming from its acoustic-based design: **1) Frequency discrepancy:** Commodity BLE chips need a 10 MHz to 50 MHz reference clock, but the tag receives only an 18 kHz to 20 kHz audio tone—three orders of magnitude lower. AcouBLE has to up-convert the frequency by almost 1000×–2000× without using active, power-hungry frequency synthesizers or oscillators. **2) Precise Audio-BLE alignment:** The smartphone must time each audio burst so it arrives at the tag immediately before the BLE connection event, allowing it to generate the clock in time. This is difficult because (a) to satisfy requirement ❶, AcouBLE cannot access the smartphone's BLE chip's internal timing, and (b) even with this information, the vast difference in propagation speed between RF (3×10^8 m/s) and acoustic (343 m/s) signals makes it nontrivial to calculate the precise transmit time without knowing the smartphone-tag distance.

To address the first challenge, we introduce a novel triangular-wave thresholding method that digitally doubles the incoming audio frequency with ultra-low jitter. Chaining multiple such stages creates a cascaded digital upconversion pipeline that rapidly boosts the signal into the MHz regime. We then exploit analog harmonic conversion at higher frequencies, where analog stages are far more power-efficient. This digital-analog synergy yields a compounding upconversion factor, achieving MHz clocks from kHz audio at only µW-level power consumption.

To address the second challenge, we introduce a lightweight feedback protocol that learns both the BLE connection interval and the precise timing offset directly from the tag, avoiding any modifications to the BLE stack or privileged access on the smartphone. Within just two connection events, the smartphone adapts its audio burst schedule so subsequent bursts align perfectly with each connection event.

Finally, to demonstrate the feasibility of our design we develop a full prototype of AcouBLE using custom PCBs (printed circuit board), off-the-shelf components, and commodity smartphones.

To summarize, we make the following key contributions:

- We present AcouBLE, the first system to dramatically cut a BLE tag's RF energy consumption through clock-seeding, a new paradigm that shifts high-frequency clock generation from the energy-constrained tag to the smartphone while preserving BLE's standard operation (§ III).
- On the tag, we introduce a hybrid digital-analog pipeline that converts a kHz audio tone into a stable 32 MHz clock using six triangular-wave frequency doublers and three low-power harmonic stages, achieving a compounding factor of 1728 at µW-level power consumption (§ IV).
- On the smartphone, we introduce a lightweight feedback protocol that learns both the BLE connection interval and the precise timing offset necessary for transmitting the audio signal without any modifications to the BLE stack or privileged access (§ V).
- We prototype AcouBLE on custom PCBs and commodity smartphones (§ VI), demonstrating up to 5.9× energy savings (§ VII) while preserving one-to-one deployment, backward compatibility, and full RF communication performance—successfully tackling the challenge outlined in our problem statement.

II. Clock-seeding BLE: Opportunities and Challenges

In this section, we revisit a fundamental question: *can the energy required to transmit and receive BLE packets in a connection event be significantly reduced—without modifying*

TABLE I
REF. CLOCK WARM-UP OVERHEAD, AS PERCENTAGE, FOR VARIOUS TX DATA SIZES: 10, 20, 50 AND 100 BYTES.

Payload size	TX Energy	Ref. Clock warm-up overhead
10 bytes	0.60 µJ	85.0%
20 bytes	0.90 µJ	83.4%
50 bytes	1.81 µJ	78.8%
100 bytes	3.29 µJ	72.3%

the BLE protocol, RF waveform, or baseband behavior—by rethinking the underlying chip operation?

To answer this question, we begin by carefully analyzing the energy breakdown of a commodity BLE in Connected mode (§ II-A). The findings of this analysis pave the way for clock-seeding (§ II-B), a new paradigm for substantially reducing the energy footprint of RF data transmissions.

A. BLE Chip's Energy Footprint in Connected Mode

In this section, we dissect the energy breakdown of an industry-leading commodity BLE chip, Nordic's NRF5340 [19], in peripheral connected mode during a connection event for different packet sizes. To accomplish this, we use Nordic's Power Profiler (PPK2) [20], a real-time nA-to-A current-profiling tool with µs resolution. As illustrated in Fig. 1, this cycle involves four major subsystems:

1) **Reference Clock Generator:** Commodity BLE chips require a stable MHz clock (typically 16/24/32 MHz) for CPU timing and RF tasks such as PLL and modulation [21], [19]. Their internal high-frequency oscillators (HFXOs) consume significant energy while converging to a low-jitter output. For the nRF5340, our measurements show the 32 MHz HFXO needs ≈4.4 ms warm-up with 1.62 mA current at 1.8 V, consuming 12.96 µJ. As Table I shows, this single step alone accounts for 72.3–85% of the total energy per BLE connection event.
2) **RF Synthesizer (PLL):** This block generates the 2.4 GHz carrier, waking up after the HFXO and locking to its stabilized MHz reference. As shown in Fig. 1, it accounts for most of the 0.33 µJ spent during the 80 µs *start radio* phase, primarily for RF PLL warm-up.
3) **RF Front-End:** Each connection event begins with the peripheral receiving a *ping* from the central device, followed by transmitting a *response*. Since synchronization is achieved during connection setup, the chip schedules wake-up so the radio start phase aligns with the ping's arrival, then immediately switches from RX to TX. During this interval, the RF front-end—modulator/demodulator, mixers, and amplifiers—is the dominant energy consumer. As shown in Fig. 1, it consumes 0.46 µJ for RX, 0.48 µJ for RX–TX switching (which takes ≈120 µs), and 0.6 µJ for TX of a packet with 10 bytes of payload at −10 dBm. In total, the RF front-end contributes ≈1.54 µJ, or 10% of the overall energy footprint.
4) **Digital baseband:** Active throughout the connection for packet synthesis and timing, this block consumes little energy. Only in the *post-processing* phase—updating buffers and decoding packets—does it dominate briefly, costing 0.41 µJ (2.8% of total energy).

Key finding: Our analysis reveals a clear bottleneck: the MHz clock generator dominates per-event energy in BLE connected mode, consuming up to 85% of total energy for small 10-byte payloads common in wearables, and still over 72% for 100-byte payloads common in high-throughput applications like audio streaming (Table I).

This finding paves the way for clock-seeding, a new paradigm for reducing the energy footprint of a commodity BLE chipset.

B. Clock-seeding: A new paradigm

The previous analysis motivates *clock-seeding*, a new paradigm for reducing BLE energy consumption. Unlike backscatter [17], [9], which offloads the entire RF front-end, clock-seeding proposes to offload only the reference MHz clock generation from tag to central. By supplying a crafted signal, the tag can reconstruct this clock at sub-µJ energy cost, sharply lowering per-event energy consumption.

Clock-seeding BLE is technically feasible because many off-the-shelf BLE chips support two clocking modes: (i) generating the MHz reference internally using a quartz crystal (XTAL) with the high-precision oscillator (HFXO); or (ii) taking an external MHz clock via the XTAL1 pin.

To show the benefit, we profile the NRF5340 when fed by an external 32 MHz clock from a Rigol DG1062Z precision signal generator. Here, the HFXO warm-up shrinks to 148 µs with 864 µA draw, i.e., 0.23 µJ, a 29.6× reduction vs. the internal oscillator. The overall energy consumed by the externally clocked BLE chip for sending the same data packet drops to ≈2.21 µJ, 6.8× lower than in internal mode. Importantly, clock-seeding retains full BLE functionality, enabling standard data exchange.

C. Clock-seeding: Core system challenge

Building a BLE solution around clock-seeding to address the fundamental problem statement (§ I) of this work raises the immediate question of what clock-seeding signal to use. It is a core design challenge shaped by requirements ❶ and ❷. ❶ (Commodity-like deployment): The signal must run during every BLE connection event while the central's BLE radio is actively transmitting and receiving. As a result, it cannot share spectrum or hardware with BLE, ruling out options like Wi-Fi. ❷ (Coexistence and backward compatibility): The signal source must not depend on interfaces that restrict third-party control, such as LTE, whose behavior cannot be modified by user applications.

III. OVERVIEW OF ACOUBLE

We present AcouBLE, a smartphone-based system that enables acoustic-aided clock seeding for BLE tags (Fig. 2). Designed to address the core system challenge described in Section II-C, AcouBLE makes two key design choices that together allow significantly reducing the energy consumption of a BLE tag while satisfying requirements ❶ and ❷.

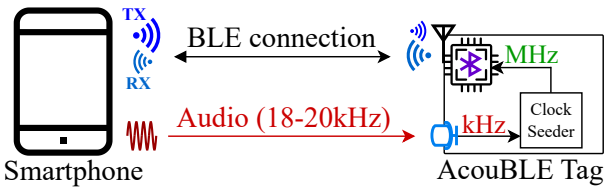

Fig. 2. AcouBLE communication system.

(1) Leveraging the smartphone as the clock source. Modern smartphones are the natural hub for ubiquitous computing: they integrate BLE radios, multiple sensing interfaces, and ample energy (relative to a tag). By offloading clock generation to the smartphone, AcouBLE transforms it into both the BLE central and the enabler for clock seeding.

(2) Using acoustics as clock-seeding signal. Requirements ❶ and ❷ exclude RF-based choices typically found on smartphones, like WiFi, 5G, UWB. Faced with this challenge, AcouBLE delivers the clock-seeding signal over the acoustic channel. Modern smartphones come equipped with high-quality speakers and audio APIs that allow precise, system-level control over sound generation. By transmitting an inaudible, heavily duty-cycled acoustic burst in the 18–20 kHz band via the system-audio path[1], AcouBLE creates a clock-seeding signal that a) coexists seamlessly with BLE communications, b) requires no hardware changes and no privileged system access on the smartphone, and c) produces no perceivable impact on the phone's audio playback (e.g., music) and introduces no privacy risks, the burst being a fixed, data-free signal.

On the tag side, a tiny, μ-powered MEMS microphone captures the acoustic tone while rejecting other audio signals such as ambient sounds. This can be done via an ultra-low-power, pre-amplifier stage after the microphone that is tuned to selectively amplify 18 kHz to 20 kHz while rejecting other frequencies. This frequency selectivity enables reliable clock extraction even in the presence of competing sounds. The microphone feeds the received acoustic burst to the clock seeding circuit, which converts it into a clean MHz signal–32 MHz in our implementation with Nordic's nRF5340 [19]—at a fraction of milliwatt power consumption and a fraction of millisecond warm-up time, thus a small fraction of µJ energy overhead. Finally, the tag processes acoustic bursts only after establishing a secure BLE connection, and the alignment protocol (§ V) ensures that only correctly timed bursts are accepted, preventing spoofing or malicious interference.

A. Design Challenges

Realizing the AcouBLE system presents several key challenges, mainly stemming from its acoustic-based clock seeding. The challenges involve the AcouBLE tag design as well as its smartphone module.

1) AcouBLE tag design challenges:

Significant frequency discrepancy: The AcouBLE tag's BLE chip requires a 10 MHz to 50 MHz reference clock, whereas the received acoustic signal lies in the 18 kHz to 20 kHz range, nearly 1000×–2000× lower. The challenge is to up-convert this frequency without power-hungry oscillators or synthesizers, using only novel µW-level circuits with minimal warm-up time.

Preserving ultra-low phase noise: Phase noise, measured in dBc/Hz, captures timing fluctuations of the clock [22]. To ensure stable RF carrier generation, the seeded MHz clock must retain <-80 dBc/Hz at 100 kHz offset [23]. Although the audio tone originates from a low-jitter smartphone PLL (§VII-A), circuit noise, misalignment, and harmonic distortion in the µW hardware can degrade spectral purity, threatening BLE reliability.

Received audio's dynamic range: The microphone output varies strongly with distance. At <0.5 m, amplitudes >10 mV can drive voltage comparators but suffer distortion from the microphone's nonlinear amplifier. At >2 m, the signal is cleaner yet weak (<2 mV), requiring amplification. Thus, clock seeding must adapt across this wide dynamic range.

2) AcouBLE smartphone module challenges:

Precise Audio-BLE alignment: Acoustic signals in the low ultrasonic range (18 kHz to 22 kHz) can be audible to pets and even perceptible to some humans. To minimize disturbance and/or interference, AcouBLE tightly duty-cycles the acoustic signal. Specifically, audio burst are transmitted only during BLE connection events, providing the tag with the clock needed to receive and transmit packets. This design introduces a key challenge: for clock seeding to succeed, each audio burst must arrive within a sub-millisecond (500 µs in our implemented clock seeder feeding Nordic's nRF5340, § VI) window immediately preceding the BLE connection event so the tag can generate the MHz clock for its commodity BLE chip to wake up in time and receive the BLE ping. Achieving this timing is difficult for two reasons. First, as a user-space module, AcouBLE cannot access the smartphone BLE chip's internal timing, including connection schedules or intervals. Second, even if this information were available, the vast difference in propagation speed between RF (3×10^8 m/s) and acoustic (343 m/s) signals makes it nontrivial to calculate the precise transmit time without knowing the distance between smartphone and tag.

IV. AcouBLE's Tag Design

The clock seeder aims to transform $X(t) = \cos(2\pi f_{in} t)$, with $f_{in} \in$ 18 kHz to 20 kHz, into a clean, ultra-low-jitter f_o = 32 MHz clock with amplitude > 200 mV_{pp}, sufficient to drive the XTAL1 pin of commodity BLE chips (Nordic nRF5340 in our prototype). This must be achieved without active oscillators, PLLs, or frequency synthesizers, ensuring sub-ms warm-up time and sub-mW power consumption.

As discussed in § III-A, the central difficulty lies in bridging a three-order-of-magnitude frequency gap, for which no passive frequency multiplier exists. To overcome this, AcouBLE employs: (a) a divide-and-conquer architecture to manage complexity, (b) a hybrid multiplier design combining two

[1]It remains active for third-party apps even when the phone is UI-muted.

limited but complementary approaches, and (c) a novel digital frequency-doubling technique.

A. A divide-and-conquer architecture

First, we observe that it is impractical to achieve the ultimate frequency multiplication factor $N = f_{out}/f_{in}$, which is an integer in the range $[1600, 1777]$, using only a single passive frequency multiplier – the multiplication factor of a single multiplier typically cannot exceed a single-digit integer value [24]. To overcome this challenge, AcouBLE breaks down the problem by factorizing N into the multiplication of single-digit integer numbers, i.e., $N = N_1 \times N_2 \times \ldots$, where each N_i is the multiplication factor of a single frequency-multiplier stage. Based on this, we introduce a clock seeder architecture consisting of a cascaded chain of frequency multiplier stages, with N_i being the multiplication factor of the i-th stage.

B. A synergistic frequency-multiplying approach

While the cascaded chain of frequency multiplier stages introduced in the previous section simplifies the complexity of the acoustic-based clock-seeding, AcouBLE still needs single-stage frequency multipliers capable of working both in the sub-MHz and MHz range. To tackle this challenge, we leverage the two prevailing design choices:

• *Digital delay-based:* This approach relies on generating a portion of the full cycle, e.g., one quarter, and creating multiple copies with different delays. Higher frequencies are then achieved by XORing these copies [25]. The key challenge here is maintaining symmetry between the signal copies to avoid timing misalignment; any mismatch in delay paths can lead to frequency distortion or increased jitter, which severely impacts the quality of the synthesized clock. To address this challenge, we introduce a novel triangular-based frequency doubler in § IV-C that enables a better control of the jitter by fine-tuning the threshold voltages directly rather than timing alignments. Nevertheless, a µW-level triangular frequency doubler does not work properly in the MHz range due to insufficient switching speed.

• *Analog harmonic-based:* Here, the input sinusoid $V_{in}(t) = \cos(2\pi f t)$ drives a nonlinear device (e.g., a transistor), producing multiple harmonics: $V_{out}(t) = \sum_i \alpha_i \cos(2\pi i f t)$. A bandpass filter (BPF) then selects only the M-th harmonic to yield frequency $M \times f$. While practical in the MHz range, this becomes infeasible in the sub-MHz range. There, adjacent harmonics (e.g., $M-1$, $M+1$) are only f apart—sometimes as low as 18–20 kHz, requiring filters with kHz-level bandwidths. Such designs introduce large delay and demand impractically large passive components (1–10 mH inductors, 10–100 mF capacitors) with high quality factors [26], [24], making them unsuitable for low-frequency operation.

Based on these observations, AcouBLE employs a hybrid approach: (a) K stages of our triangular frequency doubler handle sub-MHz frequencies, starting from f_{in} and doubling until 1 MHz (the limit of reliable comparator switching). Proper operation under varying audio levels is ensured by the adaptive duty-cycle evener (§IV-C2). (b) Above 1 MHz, up to 32 MHz, AcouBLE transitions to analog harmonic conversion stages (§IV-D), which are feasible at these frequencies using practical LC filter values.

Algorithm 1 Frequency Conversion Factorization

Require: Output frequency $f_o = 16\,\text{MHz}$; input range $(f_{\min}, f_{\max}) = (18\,\text{kHz}, 19.5\,\text{kHz})$; harmonics $\mathcal{H} = \{3, 5\}$
Ensure: Smallest valid f_i and corresponding chain
Initialize candidate set $\mathcal{C} \leftarrow \emptyset$
for $K \leftarrow 4$ to 7 **do**
  $M_{\text{base}} \leftarrow 2^K$
  for all combinations $H = H_1 \cdot H_2 \cdots H_n$ with $H_i \in \mathcal{H}$ **do**
    $M \leftarrow M_{\text{base}} \cdot H$
    $f_i \leftarrow f_o / M$
    if $f_{\min} < f_i < f_{\max}$ **then**
      Add $E = (f_i, K, n, M)$ to $\mathcal{C}$
return $E \in \mathcal{C}$ with smallest f_i (tie-break by smallest n)

Determining the cascaded chain parameters: We seek an input audio frequency $f_i \in (f_{\min} = 18\,\text{kHz}, f_{\max} = 19.5\,\text{kHz})$[2] such that the target clock frequency $f_o = 32\,\text{MHz}$ can be synthesized using a chain of digital frequency doublers and analog harmonic multipliers. The required relation is:

$$\frac{f_o}{f_i} = 2^K \cdot \prod_{i=1}^{n} H_i, \quad K \in \{4, 5, 6, 7\},\ H_i \in \{3, 5\}.$$

The bounds on K arise because harmonic multipliers are unreliable below $2^4 \times f_{\min}$, while frequency doublers fail above $2^7 \times f_{\max}$. Among all valid solutions, we choose the one with the smallest f_i; for ties, we minimize the number of harmonic stages n to reduce power and complexity. The constrained factorization is solved using Algorithm 1, which identifies the practical design point: $f_i = \frac{1}{54\mu s} \approx 18.519\,\text{kHz}$, $K = 6$, $n = 3$, with $H_1 = H_2 = H_3 = 3$. Fig. 3 illustrates the resulting clock seeder architecture.

Note that this design strategy enables a fixed-ratio frequency multiplication factor to derive the precise 32 MHz reference with minimal circuit complexity. Thus, in our current implementation of AcouBLE's clock seeder, the frequency of the acoustic burst needs to be constant. We leave the exploration of frequency-agile architectures, which can enable frequency hopping across the inaudible (18–20 kHz) band to achieve greater robustness against ambient ultrasonic interference, to future work.

C. From KHz to 1 MHz: A Novel Frequency Doubling Architecture

Challenges with XoR-based doubler's Edge Asymmetry: In the XOR-based frequency doubling method, the input square wave is delayed by a quarter-cycle $T/4$ and XORed with its original version. Ideally, this yields evenly spaced quarter-cycle pulses and a clean output at twice the input frequency. In practice, however, this is difficult—especially in ultra-low-power designs—because digital signals exhibit rising/falling edge asymmetry of 20–70% depending on I/O standard, drive strength, and load [27]. The odd quarter-cycle

[2]Beyond 19.5 kHz, the efficiency of the speaker drops significantly.

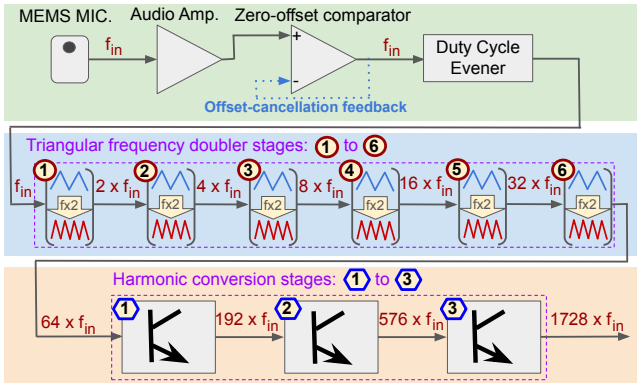

Fig. 3. The high-level architecture of AcouBLE's acoustic-based clock seeder.

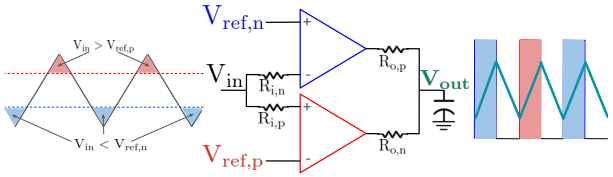

Fig. 4. Triangular frequency doubler stage.

pulse spans the rising edges of the original and delayed signals, while the even one spans their falling edges, making them inherently unbalanced.

1) Triangular Frequency Doubler: To overcome edge-induced asymmetry, we propose a novel frequency doubling method based on two architectural choices:

Design Choice 1 – Triangle Wave Conversion: Instead of XORing delayed edges, we first convert the input square wave into a triangular waveform using a passive RC low-pass filter. Unlike digital edges, triangular slopes depend only on duty cycle, making them immune to rising/falling edge mismatches. Note that the doublers do not rely on voltage gain; rather, their full-swing outputs are shaped into triangular waves by these low-pass filters.

Design Choice 2 – Quarter-Cycle Extraction via Thresholding: From this triangular wave, both even and odd quarter-cycle pulses are extracted using two comparators with high/low reference voltages. Each crossing event corresponds to a precise timing location, enabling clean pulse generation.

Threshold Selection: As shown in Fig. 4, the comparators toggle when the input crosses above $V_{\text{ref,p}} = \frac{3}{4}V_H + \frac{1}{4}V_L$ or below $V_{\text{ref,n}} = \frac{1}{4}V_H + \frac{3}{4}V_L$, where V_H and V_L are the high and low peaks of the triangular waveform (known from prior stage parameters). Each comparator is ON for 25% of the cycle, yielding pulses spaced half a cycle apart and an output frequency $f_{\text{out}} = 2f$.

Reference Voltage Implementation: Since the theoretical values of V_H and V_L are known a priori (except in the first stage), reference voltages can be precomputed and implemented via resistor-based voltage dividers. Once implemented, these thresholds remain stable since all stages share a fixed-voltage regulator. To suppress even minimal unwanted harmonics and ensure full symmetry required by later stages, thresholds pin voltages are first tuned by sweeping DAC-generated voltages in 25 µV steps. This is a one-time per-stage calibration before hard-wiring using resistor dividers.

Circuit imperfections and remedies: Although the design is inherently symmetrical, practical non-idealities can degrade output quality: • **Propagation delay imbalance:** Manufacturing variations cause unequal comparator delays, misaligning quarter-cycle pulses. We correct this by inserting a fixed RC delay at the faster comparator's input. • **Voltage offset mismatch:** Small input offsets lead to asymmetric switching. We mitigate this by fine-tuning the threshold voltages. • **Output impedance imbalance:** Driver mismatches, resistor tolerances, or load capacitance variations distort pulse shapes. We restore balance by adjusting the output load resistors.

2) Adaptive Duty-Cycle Evener: To enable acoustic-aided clock seeding, the MEMS microphone output must first be converted into a triangular waveform. In theory, using a voltage comparator with zero offset and setting its threshold to the audio's DC average would yield a 50% duty-cycled square wave and hence a symmetric triangle wave. In practice, however, nonlinear distortion in the microphone results in asymmetry between rising and falling half-cycles of the resulting triangular wave. This imbalance violates the input requirement of the first frequency doubler and cascades errors across subsequent stages. While techniques such as adaptive gain control (AGC) [28] can effectively mitigate dynamic range and temporal amplitude fluctuations in the input audio signal, they do not fully address the duty-cycle asymmetry introduced across different distances. In §VII-A, we characterize this asymmetric duty-cycle effect by analyzing the spurious inter-modulation components that grow across stages and ultimately corrupt the signal for downstream processing.

Our solution: We introduce a dedicated *duty-cycle evener* stage that adaptively corrects the asymmetry before feeding the first doubling stage of the clock seeder, Fig. 5. The objective is to achieve precisely 50% duty cycle regardless of the dynamic changes in the received audio amplitude induced by smartphone-tag distance. Our duty-cycle (DuC) evener design is based on this key observation: While asymmetric, the audio-induced triangular wave can be characterized by 3 parameters: 1. positive peak V_H, 2. negative peak V_L, and 3. duty cycle D, i.e., the fraction of the period T during which the waveform rises. To remove asymmetry, we feed this waveform into a comparator with an adaptive threshold $V_{\text{thr}} = (V_H + V_L)/2$. Geometrically, this bisects the triangle into equal-area halves with identical base widths, ensuring a precise 50% duty cycle at the comparator output. While V_H is obtained directly using a low-power op-amp peak detector, measuring V_L is harder. Instead, we compute it analytically: during the falling edge, the waveform follows $V_f(t) = V_H \cdot (1 - t/\tau)$, where τ is the known time constant of the RC filter generating the triangular wave. Since $V_L = V_f((1-D)T)$, we have $V_L = V_H \cdot \left(1 - \frac{(1-D)T}{\tau}\right)$. The DC average relates linearly to the duty cycle as $V_{\text{avg}} = D \cdot V_{cc}$, yielding $V_L = V_H \cdot \left(1 - \frac{T}{\tau} + \frac{T}{\tau \cdot V_{cc}} \cdot V_{\text{avg}}\right)$. Thus, the adaptive threshold becomes $V_{\text{thr}} = \frac{V_H + V_L}{2} = V_H \cdot (\alpha \cdot V_{\text{avg}} + \beta)$,

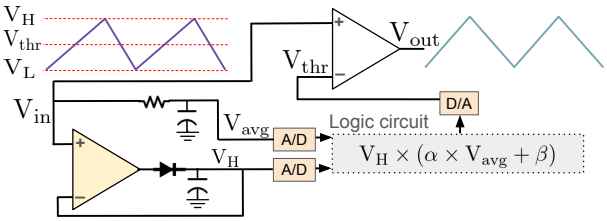

Fig. 5. Adaptive duty-cycle evener.

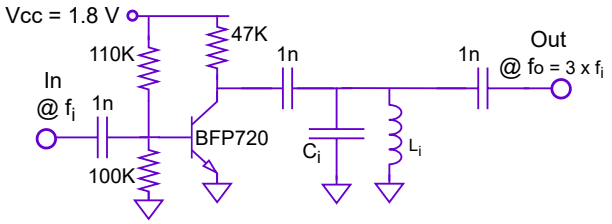

Fig. 6. MHz (3rd) harmonic conversion stage.

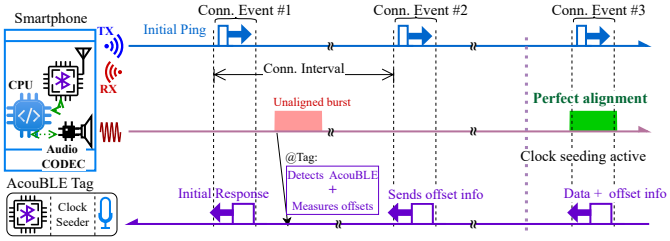

Fig. 7. AcouBLE's protocol for acoustic-BLE timing alignment.

where $\alpha = 1 - \frac{T}{2\tau}$, $\beta = \frac{T}{2\tau \cdot V_{cc}}$.

D. MHz Regime: Harmonic Conversion

In the input frequency range above 1 MHz, ultra-low-power (µW) comparators fail to support clean frequency doubling due to limited switching speed. Hence, beyond this point we transition to harmonic-based signal generation, which becomes practical as passive harmonic extraction with LC bandpass filters is far less challenging at MHz frequencies. Resonant filters with realizable capacitor and inductor values can thus be tuned to specific harmonics.

To upconvert from $64 \times \frac{1}{54\,\mu s} \approx 1.185$ MHz to 32 MHz, we use three cascaded 1:3 harmonic conversion stages. As shown in Fig. 6, each stage employs a BJT in common-emitter mode, biased into nonlinearity to generate harmonics, with the LC filter selecting only the 3rd harmonic. The first stage outputs 3.55 MHz, the second 10.65 MHz, and the third 32 MHz. While all stages share the same topology, LC values differ by frequency: $\{L_1 = 3\,\mu H, C_1 = 680\,pF\}$ at stage one, $\{L_2 = 330\,nH, C_2 = 680\,pF\}$ at stage two, and $\{L_3 = 160\,nH, C_3 = 150\,pF\}$ at stage three. Each stage achieves >40 dB suppression of unwanted harmonics with a single-order BPF, while maintaining <10 µs startup time. Finally, large collector resistors (47 kΩ) keep overall power consumption in the µW regime (see §VII-B).

V. AcouBLE's Smartphone Module

In this section, we introduce AcouBLE's smartphone module. As described in § III-A2, its key challenge is the precise alignment of the heavily duty-cycled audio signal transmission with BLE connection events.

A. Acoustic alignment with BLE events

In BLE connected mode, the central (smartphone) and peripheral (tag) exchange packets periodically during connection events. Each event begins with the central sending a Ping packet—even if no data needs to be transmitted—after which the peripheral may reply before both devices return to sleep until the next event. The connection interval, negotiated during discovery, can range from 7.5 ms to 4000 ms [29].

As AcouBLE tightly duty-cycles the acoustic signal to minimize interference, each burst must be precisely aligned with the BLE connection event. In particular, the burst must arrive immediately before the event so the tag can seed its clock (§ IV) and wake up in time to receive the BLE ping and send its response. Achieving this timing is challenging for two reasons: First, the connection interval and start time of each event are known to the smartphone's BLE chip but not exposed to user space. Second, even with this information, propagation delay asymmetry between RF and acoustic signals makes it impossible to predict when to start the audio transmission without knowing the smartphone-tag distance.

Moreover, to satisfy requirement ❷ (backward compatibility), any solution must ensure that tags and smartphones behave exactly like commodity BLE devices unless both ends explicitly recognize each other as AcouBLE-enabled.

AcouBLE overcomes these challenges with a lightweight, backward-compatible feedback protocol (Fig. 7) that leaves the smartphone BLE stack unmodified while leveraging full control over the tag firmware. The protocol proceeds in three phases:

1) **Discovery:** The tag initially operates as a standard BLE device (using its internal MHz oscillator), transmitting a packet in the first connection event.
2) **Initial burst & recognition:** Upon receiving this packet, the smartphone's AcouBLE module sends an unaligned audio burst. This first burst serves two purposes: it confirms to the tag that the central is AcouBLE-enabled and provides a reference point for timing calibration.
3) **Feedback & alignment:** The tag measures the offset between the burst's actual and ideal arrival times and returns this offset, along with the connection interval, in the next BLE packet. To avoid modifying the format of a BLE packet, this information is inserted in the payload and removed at reception by the smartphone module. Using this feedback, the smartphone aligns subsequent audio bursts precisely so that, by the third connection event, bursts arrive exactly when needed—enabling reliable clock seeding without stack modifications or privileged access.

B. Multi-tag operation

While described for a single tag, the acoustic alignment protocol naturally extends to multi-tag deployments by leveraging BLE connected-mode scheduling. In BLE, the central

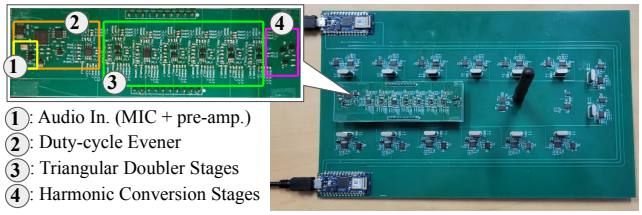

Fig. 8. PCB prototype of AcouBLE's clock seeder.

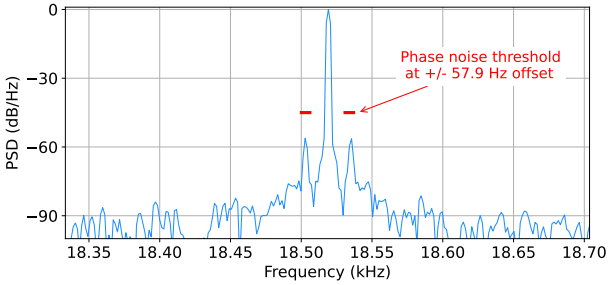

Fig. 9. Phase noise of the incoming acoustic signal.

assigns each peripheral a distinct connection interval and anchor point[3], ensuring that connection events do not overlap in time. AcouBLE runs the alignment protocol independently with each tag during its own connection events. As a result, each tag processes the acoustic signal only during its scheduled events and ignores bursts associated with other connections. Multiple AcouBLE tags can therefore operate concurrently within the same acoustic environment without contention. Scalability is ultimately bounded by the central's supported number of concurrent BLE connections, rather than by acoustic interference.

VI. Implementation

A. AcouBLE Tag

Fig. 8 shows the prototype clock seeder, designed as a daughter-board on a custom motherboard. Key components include:

Microphone and pre-amplifier: TDK ICS-40214 [30] MEMS microphone (optimized for 18 kHz to 19.5 kHz) with a pre-amp built using TI OPA2333P [31] ultra-low-power op-amp, providing 20 dB gain at 18.518 kHz. The same op-amp supports peak detection in the duty-cycle evener (Fig. 5).

Clock seeder: (i) Duty-cycle evener: AMD XC2C32A CPLD [32] with TI ADS7042 [33] 12-bit ADC, DAC7562 [34] 12-bit DAC, and TLV7011 [35] comparator. (ii) Triangular doubler chain: TI TLV7012 dual comparators [35]. (iii) Harmonic conversion: Infineon BFP720 BJTs [36] (Fig. 6).

BLE subsystem: Clock seeder connects to a custom motherboard hosting Nordic nRF5340 BLE SoC, which can select between an onboard 32 MHz crystal and the external clock.

Fine-tuning: High-precision thresholds are set using TI DAC1220 [37] 20-bit DACs, controlled by two Arduino Nano 33 IoT boards [38]. These DACs sweep comparator voltages for calibration, later fixed via resistor dividers. Also, we use a RIGOL DHO814 oscilloscope to analyze microphone input, intermediate triangular stages, harmonic outputs, and the final 32 MHz clock.

B. AcouBLE smartphone module

The central device is a Samsung Galaxy S21+ running Android 14. The AcouBLE smartphone module is implemented as a custom Android Studio app. It controls acoustic bursts' timing and duty cycling while managing BLE events. Although the operating system introduces some scheduling variability, our measurements show that OS-induced jitter remains below 40 µs per connection event. In our implementation, this variability is absorbed by a small timing margin in the acoustic burst schedule. Moreover, because bursts are triggered relative to the BLE anchor point, the protocol re-synchronizes at every interval, preventing jitter accumulation.

[3]The start time of a connection event, when the central sends the first packet, recurring every connection interval.

VII. Evaluation

A. Justification of design choices

Incoming audio phase noise: The clock seeder's performance highly depends on the quality of the audio tone received by the MEMS microphone. The final 32 MHz output must satisfy $-80\,\mathrm{dBc/Hz}$ at a 100 kHz offset [23], i.e., the noise power density at $32\,\mathrm{MHz}\pm100\,\mathrm{kHz}$ must not exceed $-80\,\frac{\mathrm{dB}}{\mathrm{Hz}}$. For the input audio tone $f_{in} = \frac{32}{64\times27}\,\mathrm{MHz} \approx 18.518\,\mathrm{kHz}$, this translates to requiring that noise at $f_{in}\pm\frac{100}{64\times27}\,\mathrm{kHz} \approx f_{in}\pm57.9\,\mathrm{Hz}$ not exceed $-80\,\mathrm{dB}/(\frac{1}{64\times27}\,\mathrm{Hz})$, or equivalently $-47.6\,\mathrm{dB/Hz}$.

Fig. 9 shows the spectral power density (PSD) of 1 s-long audio samples at 1 ns resolution when the distance between the MEMS microphone and the smartphone's speaker is 2 meters; In this case, the noise power density at ±57.9 Hz offset remains ≈ 13 dB/Hz below the required threshold, confirming that the audio signal is inherently suitable as input to a properly designed clock seeder.

The need for fine-tuning the triangular frequency doublers: Fig. 10 shows the signals generated by our implemented triangular frequency doubler chain, consisting of 6 cascaded stages, before fine-tuning the component values. The chain is driven by 50% duty-cycled input at $f_i \approx 18.518\,\mathrm{kHz}$. It is observed that while we accurately set the threshold voltages according to formulas in § IV-C, circuit imperfections cause significant imbalance in the frequency-doubled triangular waves, thus causing jitter that gets compounded across stages. At the final output, the total power of unwanted terms within ±3.7 kHz of the target frequency is only ≈27 dB below the main signal power—rendering the output far from suitable for clock seeding. This highlights the necessity of fine-tuning resistor values at each frequency doubling stage to account for circuit and component imperfections. Fig. 11 shows the signals after fine-tuning. It is observed that fine-tuning effectively mitigates jitter by throttling the distortion terms that would otherwise propagate and amplify. At the final stage, the total

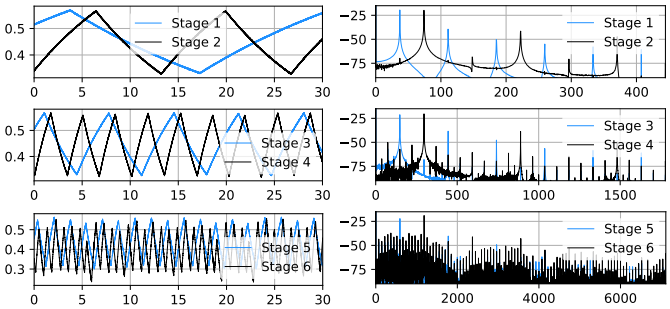

Fig. 10. Triangular doubler chain performance: without fine-tuning.

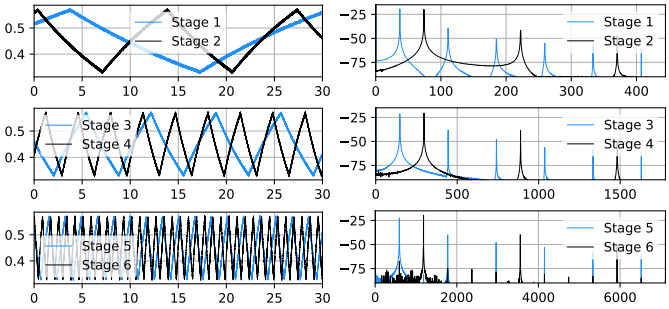

Fig. 11. Triangular doubler chain performance: with fine-tuning.

distortion power around the target frequency is reduced by over 50 dB. Also, the jitter at $\pm 3.7\,\mathrm{kHz}$ offset falls 16 dB below the required $-65\,\mathrm{dBc/Hz}$ threshold, making the final output highly robust for driving the harmonic conversion stages in the MHz regime.

The need for adaptive duty-cycle evener: As discussed in § IV-C2, even a fine-tuned triangular doubler chain cannot be directly fed with the microphone + pre-amp.'s output, due to the the imbalanced rising and falling edges inherent in that signal. Figs. 12 and 13 show this problem at two smartphone-tag distances: 50 cm and 2 m. In both cases, the comparator produces an imbalanced clock: ≈48% duty cycle at 50 cm and ≈37% at 2 m. It is observed that in both cases, similar to Fig. 10, the imbalance propagates across the chain and fills the 6th stage output with high-intensity spurious terms.

Figs. 12 and 13 also show the outputs of the stages after applying our proposed adaptive duty-cycle evener (§ IV-C2). It is observed that at both tag-smartphone distances the evener restores performance to nearly match the case when it is fed by a 50% duty-cycled input (Fig. 11). More specifically, the jitter at $\pm 3.7\,\mathrm{kHz}$ offset from 1.185 MHz is 11 dB and 13 dB below the required threshold at 50 cm and 2 m, respectively. This demonstrates the duty-cycle evener's ability to make the clock seeder robust to large position-induced dynamic changes.

B. Micro-benchmarks

Clock seeder power consumption: We evaluate the power consumption of our implemented clock seeder using Nordic's Power Profiler Kit II (PPK2) [20]. The reported numbers already include all mixed-signal overhead, including the duty-cycle evener's ADC, DAC, and CPLD. Each block's current

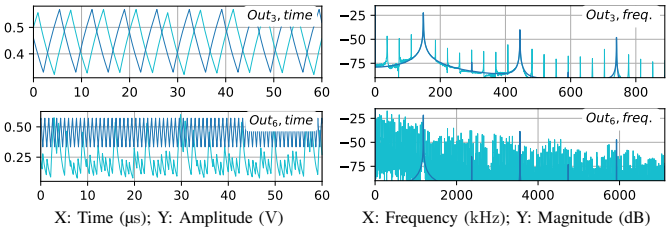

Fig. 12. Clock seeder's stage 3 and 6 outputs at 50 cm phone-tag distance. Light blue: without DuC. evener. Dark blue: with DuC. evener.

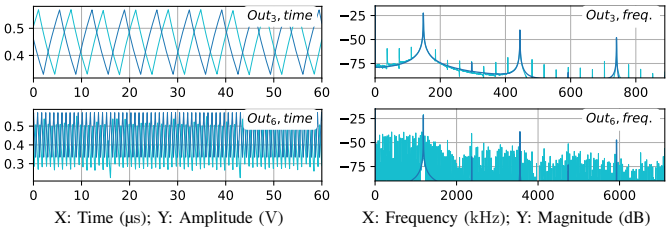

Fig. 13. Clock seeder's stage 3 and 6 outputs at 2 m phone-tag distance. Same coloring as Fig. 12.

was measured while the full pipeline was active. **Audio in:** the MEMS microphone and pre-amp draw 26.9 µA (48.4 µW). **Duty-cycle evener:** the digital logic consumes 1.8 µA and the comparator 3.9 µA, totaling 10.3 µW. The evener's low-µA draw, despite using these components, is due to the <20 kHz operating frequency. **Sub-MHz doublers:** each stage draws 3.1–3.6 µA, totaling 19.9 µA (35.8 µW). **MHz harmonic stages:** the three cascaded BJTs draw 34.7 µA, 47.1 µA, and 59.9 µA (255 µW). Overall, the clock seeder consumes 349.5 µW while supplying a stable 32 MHz clock.

ASIC Integration: Integrating the clock-seeder circuit into an ASIC/SoC is feasible and requires minimal silicon. The doublers, harmonic multipliers, and the evener map cleanly to modern low-power CMOS blocks: <10 µm per stage and power overhead in the low-µW, far below crystal oscillators.

Energy footprint comparison: Table II compares per-event energy of commodity BLE and AcouBLE tags across TX payload sizes (Fig. 1 details the commodity BLE consumption). Both are built on Nordic's nRF5340. **Commodity BLE tag:** Working with its internal crystal-based MHz clock generator, the commodity tag's BLE chip requires 12.96 µJ to wake up (Fig. 1). After waking up, it consumes 1.98 µJ to 4.68 µJ for the rest of the connection event, depending on the TX payload size. Thus, the commodity BLE tag consumes 14.94 µJ to 17.64 µJ (Table II) in total per connection event.

AcouBLE tag: The clock seeder takes ≈112 µs to warm up, resulting in 39.2 nJ overhead. It must also remain active throughout the BLE chip's wake-up and connection event. Depending on the TX payload size, this lasts between 810 µs and 1.17 ms, corresponding to an additional energy cost of 283.1 nJ to 408.9 nJ. Overall, the clock seeder consumes 322.3 nJ to 448.1 nJ per event (Table II). The externally clocked BLE chip (nRF5340) requires only 147.8 µs to wake up while drawing 864.3 µA, resulting in 230 nJ energy overhead. After waking up, it consumes 1.98 µJ to 4.68 µJ depend-

TABLE II
PER-EVENT ENERGY: ACOUBLE VS. COMMODITY BLE TAG.

TX payload size (Bytes) →	10	20	50	100
AcouBLE's clock seeder	322.3 nJ	336.2 nJ	378.2 nJ	448.1 nJ
AcouBLE's BLE chip	2.21 µJ	2.51 µJ	3.41 µJ	4.91 µJ
AcouBLE's total	2.53 µJ	2.85 µJ	3.79 µJ	5.36 µJ
Commodity BLE tag's total	14.94 µJ	15.24 nJ	16.14 nJ	17.64 nJ
Commodity tag/AcouBLE	5.9×	5.3×	4.3×	3.3×

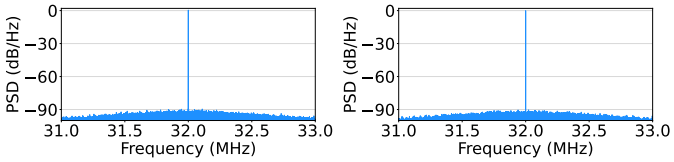

Fig. 14. Power spectral density of the clock seeder's final 32 MHz output. Left: 50 cm phone-tag distance. Right: 2 m phone-tag distance.

ing on the TX payload byte size to complete a connection event. Thus, the AcouBLE tag's total per-event energy is 2.53 µJ to 5.36 µJ, which is 3.3×–5.9× lower than a commodity BLE tag (Table II).

32 MHz clock phase noise: Fig. 14 shows the spectral density of the seeded 32 MHz clock at 50 cm and 2 m smartphone-tag distances. In both cases, noise power at ±100 kHz offset remains more than 10 dB below the required −80 dBc/Hz, demonstrating that the seeded clock achieves the fidelity required for nRF5340 RF communication.

C. Macro-level benchmarks

Experimental setup: We evaluate the end-to-end performance of AcouBLE using a Galaxy S21 smartphone with our custom app for audio transmission and timing/frequency alignment and compare it with that of a commodity BLE chip. To ensure inaudibility, the phone transmits audio bursts at ≈50 mV speaker output. Experiments are conducted in three environments (Fig. 15): (i) a 5 m×6 m empty room (line-of-sight), (ii) a cluttered 3.2 m×4.1 m office (multipath), and (iii) a 1.6 m-wide corridor (confined).

Clock-seeding range: Fig. 16 shows the clock seeding performance vs. distance. Successful clock-seeding requires the received audio amplitude >1.6 mV (peak-to-peak) at the clock-seeder's input. Thanks to the audio pre-amp, the range extends to >6 m in the empty room, ≈3.3 m in the office, and >10 m in the corridor, demonstrating sufficient indoor coverage[4] for short-range applications (e.g., smart homes, wearables, logistics, etc.).

Clock-seeding range across different smartphones: To evaluate robustness across smartphone hardware, we measured the acoustic output of devices spanning different manufacturers and generations (Galaxy S9, iPhone 13, and Google Pixel 8) and compared them to our primary platform (Galaxy S21+). Using the *Spectroid* app [39], we measured the received signal strength of inaudible tones in the 18–20 kHz band. Across devices, our measurements show minimal (<2 dB) variation in RSS versus distance at both 18.5 kHz and 19.5 kHz, indicating that differences in smartphone speakers and audio pipelines do not materially affect AcouBLE's operation or deployability.

[4]At distances beyond clock-seeding range, AcouBLE tags fall back to commodity operation, i.e., using the BLE chip's internal MHz oscillator.

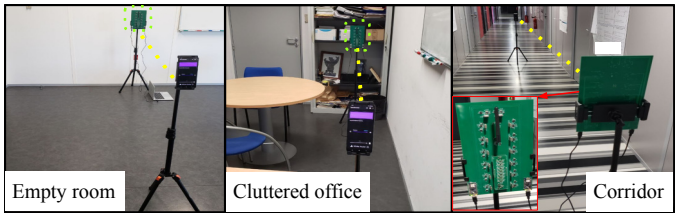

Fig. 15. Experimental setup for evaluating macro-level benchmarks.

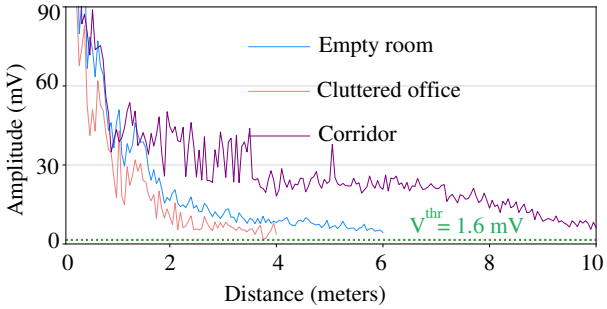

Fig. 16. Clock seeder's input audio amplitude vs. distance; with pre-amplifier.

RF performance: We next evaluate the packet delivery rate (PDR) of AcouBLE within the above seeding ranges and compare it to commodity BLE. Using an nRF52840 sniffer dongle [40] + Wireshark on a HP laptop, we monitor uplink (tag→phone) and downlink (phone→tag) reliability. Packet errors are detected by repeated ping messages (uplink) or lack of tag response (downlink). For evaluation, the connection interval is set to 20 ms, yielding 500 connection events over each 10 s run.

Static deployment: Fig. 17 shows uplink and downlink packet delivery rate vs. smartphone-tag distance for static positions in the three environments (empty room, office, corridor). In all cases AcouBLE achieves almost identical performance with commodity BLE. There is only a slight drop in the cluttered office scenario at the biggest distance, which occurs due to strong multipath impacting the audio signal.

Robustness against mobility: We repeat the packet delivery measurements with smartphone horizontally moving at 1 m/s. Fig. 18 shows that AcouBLE remains robust against mobility. We only observe a small drop (less than 20%) at the longest ranges due to weaker signals causing less reliable frequency/timing alignment from audio bursts.

Robustness against noise/interference: We evaluate packet delivery rate with a second Galaxy S9 playing (i) fan noise, (ii) recorded speech, and (iii) music, at medium volume, 2 m away and 45° off-axis. Fig. 19 shows AcouBLE remains robust to fan/speech. Even in the most challenging scenario (long-range, music) packet delivery stays above 70%.

Preserving pet-inaudibility: Finally, we analyze inaudibility at maximum phone-tag connectivity distance in the 5 m×6 m empty room. Fig. 20 shows received audio dB-SPL heatmap and contours for dog (≥50 dBSPL) and cat

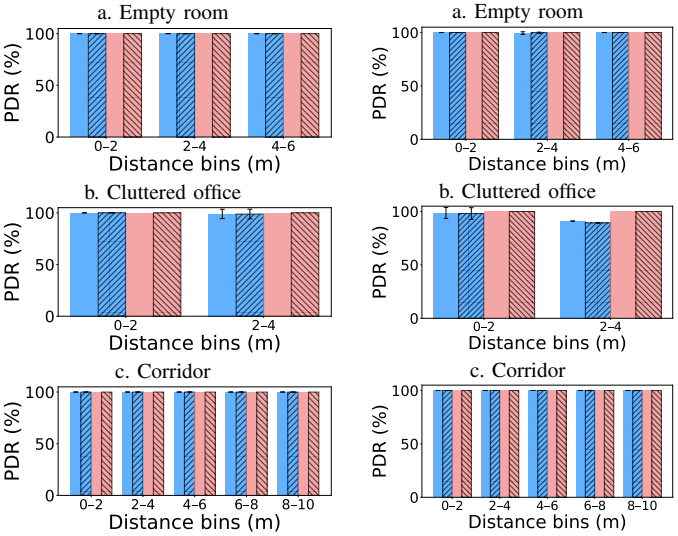

Fig. 17. PDR vs. distance, static deployment, AcouBLE in blue, commodity BLE in orange. Solid: uplink, hatched: downlink.

Fig. 18. PDR vs. distance, mobile scenario, same coloring as Fig. 17.

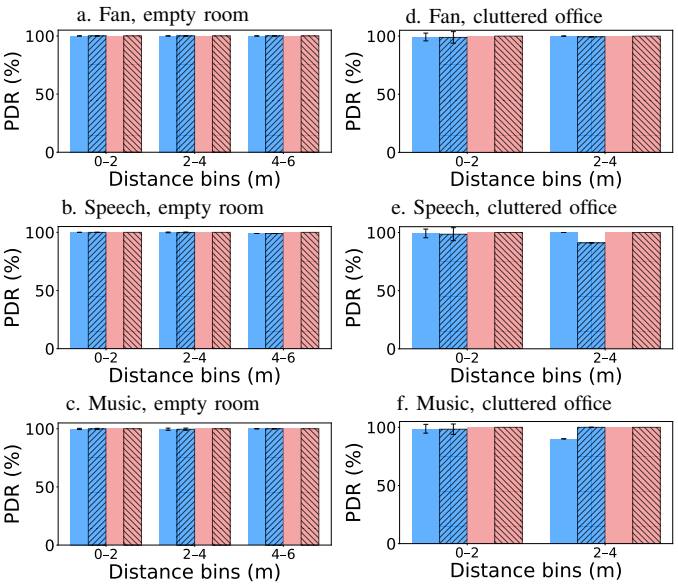

Fig. 19. PDR vs. Ambient sound, same coloring as Fig. 17.

($\geq$45 dBSPL) thresholds [41], [42]. Results show inaudibility to dogs beyond 60 cm, and to cats beyond 1.3 m.

VIII. Related Work

Energy-optimized BLE: Prior work reduces BLE energy via system- and firmware-level optimizations [5], [43] such as adaptive connection intervals, duty cycling, lightweight cryptography, and hardware co-design at the baseband and RF front-end [5], [6], [7], [8]. While effective, they remain fundamentally limited by the active RF front-end's energy cost. AcouBLE tackles this cost by offloading the high-power on-chip clock generator, enabling orthogonal and complementary savings when combined with existing optimizations.

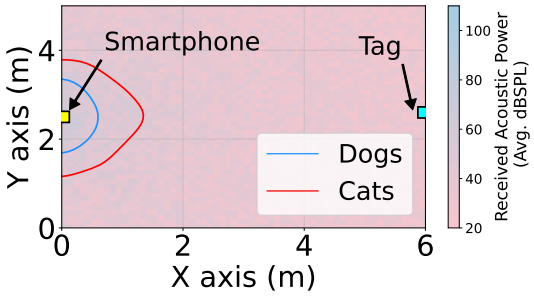

Fig. 20. The pet-audible regions around the smartphone.

Acoustic-aided IoT: Recent work leverages phone audio for communication, localization, and sensing [44], [18], [45], exploiting ubiquitous speakers and microphones at ultrasonic or near-ultrasonic bands [46]. While inspired by this, AcouBLE differs: existing systems use audio as a communication medium, whereas AcouBLE uses it only for clock seeding, avoiding the latency and robustness limits of standalone acoustic links [47].

Backscatter with commodity devices: Prior work demonstrates that WiFi, Bluetooth, or FM radios can serve as exciters for tags [13], [11], [10], [9], [48], [49], [14], [50], enabling ultra-low-power connectivity by eliminating the RF front-end entirely. While this yields unparalleled energy savings it comes at the cost of specialized infrastructure and potential network interference, limiting practical deployment [16]. In contrast, AcouBLE keeps standard BLE hardware yet slashes energy use via clock-seeding, enabling low-power, one-to-one, backward-compatible operation.

IX. Conclusion

We presented AcouBLE, the first system to cut BLE tag energy consumption by up to 5.9$\times$ while preserving one-to-one deployability, backward compatibility, and full RF performance. Enabled by clock-seeding, a new paradigm that offloads MHz clock generation to the smartphone, AcouBLE combines a lightweight timing protocol with a hybrid digital-analog upconversion pipeline to deliver ultra-low-power operation without BLE stack modifications or added infrastructure. Looking ahead, future work can extend AcouBLE to support other wireless protocols and explore adaptive techniques for improving robustness in more diverse acoustic environments.

More broadly, AcouBLE illustrates the untapped potential of cross-capability synergy in modern IoT ecosystems. We hope this inspires future designs that move beyond siloed, single-task approaches and instead harness the full suite of sensing, communication, and computation resources in today's smart devices.

Acknowledgment: We thank the anonymous reviewers and our shepherd for their insightful and constructive feedback. This work was supported in part by the CNRS Starting Grant and the ANR AGREENET grant.

REFERENCES

[1] "Edge of things (eot) for telecommunications," IBM, 2023, iBM Think Insights.

[2] IoT-Analytics, "State of iot summer 2024," 2024, forecast report, IoT-Analytics.

[3] G. Koulouras, G. P. Moustakas, G. E. Bertsias, and P. K. Papadopoulos, "Evolution of bluetooth technology: Ble in the iot ecosystem," *Sensors*, vol. 25, no. 4, p. 996, 2025.

[4] "What is the power consumption of generic bluetooth earphones," ChinaWence, 2024, blog article.

[5] J. de Winkel, H. Tang, and P. Pawełczak, "Intermittently-powered bluetooth that works," ser. ACM MobiSys '22, 2022.

[6] N. Djidi *et al.*, "Wubble: Energy efficient ble neighborhood discovery," in *Proceedings of ICT-Energy / related workshop (2023)*, 2023.

[7] C. Rup, Q. Hopp, S. Mamert, B. Turco, E. Bajic, and K. Mekki, "Experimental analysis of a bluetooth low energy wake-up radio solution," in *13th International Workshop on Service-Oriented, Holonic and Multi-Agent Manufacturing Systems (SOHOMA)*. Springer, Sep. 2023.

[8] L. Schulthess, S. Cortesi, and M. Magno, "Wakemod: A 6.9 uw wake-up radio module with -72.6 dbm sensitivity for on-demand iot," *arXiv preprint arXiv:2505.21529*, 2025.

[9] P. Zhang, C. Josephson, D. Bharadia, and S. Katti, "Freerider: Backscatter communication using commodity radios," ser. ACM CoNEXT '17, 2017.

[10] V. Liu, A. Parks, V. Talla, S. Gollakota, D. Wetherall, and J. R. Smith, "Ambient backscatter: wireless communication out of thin air," ser. ACM SIGCOMM '13, 2013.

[11] A. Wang, V. Iyer, V. Talla, J. R. Smith, and S. Gollakota, "FM backscatter: Enabling connected cities and smart fabrics," in *USENIC NSDI 17*, 2017.

[12] S. Naderiparizi, M. Hessar, V. Talla, S. Gollakota, and J. R. Smith, "Towards battery-free hd video streaming," ser. NSDI'18, 2018.

[13] M. Rostami, K. Sundaresan, E. Chai, S. Rangarajan, and D. Ganesan, "Redefining passive in backscattering with commodity devices," ser. ACM MobiCom '20, 2020.

[14] H. Dong, Y. Wu, F. Li, W. Kuang, Y. He, Q. Zhang, and W. Wang, "Passiveble: Towards fully commodity-compatible ble backscatter," in *ACM MobiCom '25, November 4-8, 2025, Hong Kong, China*, 2025.

[15] V. Talla, J. Smith, and S. Gollakota, "Advances and open problems in backscatter networking," *ACM GetMobile'21*, 2021.

[16] F. Dehbashi, A. Abedi, T. Brecht, and O. Abari, "Verification: can wifi backscatter replace RFID?" in *ACM MobiCom*, 2021, p. 97–107.

[17] V. Iyer, V. Talla, B. Kellogg, S. Gollakota, and J. Smith, "Inter-technology backscatter: Towards internet connectivity for implanted devices," ser. ACM SIGCOMM '16.

[18] S. Cao, D. Li, S. I. Lee, and J. Xiong, "Powerphone: Unleashing the acoustic sensing capability of smartphones," in *ACM MobiCom '23*, 2023.

[19] *nRF5340 SoC Product Brief*, Nordic Semiconductor.

[20] *Power Profiler Kit II (PPK2) Product Brief, Version 1.0*, Nordic Semiconductor, 2020, datasheet.

[21] *NXH3670UK Ultra-low power 2.4 GHz BLE Transceiver Data Sheet*, NXP Semiconductors, 2022, rev. 3.4, Section "Clock Generation".

[22] K. W. Yeom and Y. Roh, "An efficient cross-correlation method for a digital phase noise measurement system," *Journal of Electromagnetic Engineering and Science*, vol. 22, no. 6, pp. 675–683, 2022.

[23] https://devzone.nordicsemi.com/f/nordic-q-a/49590/using-an-external-oscillator-tcxo-with-a-nrf52-as-input-to-the-hfxo.

[24] Y. Ding, C. Wang, X. Wang, C. Huang, and B. Zhou, "A high-efficiency frequency multiplier with triangular-resistance phase interpolation," *Preprint (2025)*, 2025.

[25] G. Mazin, A. Stejskal, M. Dudka, and M. Ježek, "Non-blocking programmable delay line with minimal dead time and tens of picoseconds jitter," *arXiv preprint arXiv:2105.09183*, 2021.

[26] S. Im, K.-H. Nam, and J. Park, "A single-stage 12-times frequency multiplier for a 5g frequency synthesizer," *Journal of Electromagnetic Engineering and Science*, vol. 22, no. 3, pp. 302–308, 2022.

[27] N. Ren, Z. Fu, S. Lei, H. Liu, and S. Tian, "Methodology for digital synthesis of deterministic and random jitter generation on rising or falling edges of data pattern," *Electronics*, vol. 8, no. 12, p. 1510, 2019.

[28] D. Panchal and A. Naik, "0.25v bulk driven variable gain amplifier using gmb/idrain method for ultra-low power low voltage biomedical applications," in *2020 5th International Conference on Communication and Electronics Systems (ICCES)*, 2020, pp. 166–170.

[29] J. Tosi, M. Alfaro, M. Mernik, and D. Griol, "Performance evaluation of bluetooth low energy," *Sensors*, vol. 17, no. 4, p. 850, 2017.

[30] *ICS-40214 Analog MEMS Microphone Datasheet*, TDK InvenSense.

[31] *OPA2333P 1.8-V, μ-Power Zero-Drift CMOS Operational Amplifier Datasheet*, Texas Instruments, 2017, revision A.

[32] *DS310 XC2C32A CoolRunner-II CPLD Data Sheet*, AMD / Xilinx.

[33] *ADS7042 Ultra-Low Power, Ultra-Small Size, 12-Bit 1-MSPS SAR ADC Datasheet*, Texas Instruments, revised December 2015 (Rev. C).

[34] *DAC7562 Dual 12-Bit Buffered Voltage-Output DAC With 2.5 V Internal Reference Data Sheet*, Texas Instruments, 2015, revision E, June 2015.

[35] *TLV7011 Micro-Power Low-Voltage Comparator Data Sheet*, Texas Instruments, 2017, revision F, October 2019.

[36] *BFP720 SiGe:C NPN RF Bipolar Transistor Data Sheet*, Infineon Technologies AG, 2018, revision 2.0, September 2018.

[37] *DAC1220 20-Bit Low-Power Digital-to-Analog Converter Data Sheet*, Texas Instruments, 2009, rev. September 2009.

[38] *Arduino Nano 33 IoT Product Reference Manual (ABX00027)*, Arduino, 2025, modified: 16/09/2025.

[39] "Spectroid," https://play.google.com/store/apps/details?id=org.intoorbit.spectrum.

[40] *nRF52840 Dongle Product Brief*, Nordic Semiconductor.

[41] C. Guérineau, A. Broseghini, M. Lõoke, G. Dehesh, P. Mongillo, and L. Marinelli, "Determining hearing thresholds in dogs using the staircase method," *Veterinary Sciences*, vol. 11, no. 3, p. 68, 2024.

[42] H. Heffner and R. Heffner, "Hearing range of the domestic cat," *Hearing Research*, vol. 19, no. 1, pp. 85–88, 1985.

[43] Y. Li, B. Li, J. Lv, and W. Dong, "Bledge: Edge-centric programming for ble applications with multi-connection optimization," *ACM Trans. Sen. Netw.*, 2024.

[44] M. Rostami, A. Liu, and K. Sundaresan, "Scalable acoustic iot through composable distributed beamforming tags," in *IPSN'24*, 2024, pp. 39–50.

[45] G. C. L. Delfa *et al.*, "Survey of smartphone-based datasets for indoor localization," *Sensors*, vol. 25, no. 7, p. 2508, 2025.

[46] W. Luo, Q. Song, Z. Yan, R. Tan, and G. Lin, "Indoor smartphone slam with learned echoic location features," *Proceedings of the ACM/IEEE International Conference on Internet of Things Design and Implementation (IoTDI)*, 2022.

[47] C. Cai, R. Jin, J. Nie, J. Kang, Y. Zhang, and J. Luo, "Reliable high throughput aerial acoustic communication for mobile network," *IEEE Transactions on Vehicular Technology*, 2024.

[48] V. Talla, M. Hessar, B. Kellogg, A. Najafi, J. R. Smith, and S. Gollakota, "Lora backscatter: Enabling the vision of ubiquitous connectivity," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 2017.

[49] A. Varshney, A. Soleiman, and T. Voigt, "Tunnelscatter: Low power communication for sensor tags using tunnel diodes," ser. ACM MobiCom '19, 2019.

[50] D. Bharadia, K. R. Joshi, M. Kotaru, and S. Katti, "Backfi: High throughput wifi backscatter," ser. ACM SIGCOMM '15, 2015.