

# Acoustic Side-Channel Communications for Autonomous Drones via BLDC Motor Magnetostriction

Dakota W. Winslow

May 1, 2026

## Abstract

This report presents a hardware-minimal acoustic side-channel communication system for autonomous quadrotor drones, in which the platform’s existing brushless DC (BLDC) propulsion motors are repurposed as ultrasonic data transmitters. The approach exploits magnetostriction in the motor’s ferromagnetic structure: a high-frequency carrier signal is mathematically superimposed on the open-loop sinusoidal commutation waveform and emitted as ultrasonic acoustic radiation by the motor bell. A single MEMS microphone paired with an ESP32-S3 microcontroller serves as the receiver, demodulating Binary Frequency-Shift Keyed (BFSK) data via the Goertzel algorithm and recovering structured packets framed with CRC-16 integrity checks. Bench measurements confirm clear spectral lines through 5.6 m without propeller load, identify a strong 19 kHz acoustic resonance under load, and demonstrate end-to-end packet decoding at 1.8 m with 100% retention. The work establishes a foundational acoustic channel intended to support future drone-to-drone time-of-flight ranging in RF-denied environments.

## 1 Introduction

Autonomous quadrotor drones are increasingly deployed into operational environments where conventional radio-frequency (RF) communication is unreliable or impractical. Active warzones, high-voltage power infrastructure, and shielded industrial facilities all introduce intentional jamming, broadband electromagnetic interference, or structural attenuation that compromise the 2.4 GHz and 5.8 GHz links on which most drones depend. RF emissions are also long-range and broadcast-like, making them unsuitable for missions in which the operator wishes to remain low-observable. Acoustic communication, by contrast, attenuates rapidly in air, is intrinsically short-range and directional, and is physically separated from the electromagnetic spectrum. A drone capable of communicating acoustically — particularly in the ultrasonic band, where the signal is inaudible to humans — gains a complementary channel that survives where RF fails.

The principal challenge in deploying such a channel is hardware. Adding a dedicated ultrasonic transducer, amplifier, and microcontroller to a weight- and power-constrained drone undermines the practicality of the approach. The most elegant alternatives are those that repurpose existing hardware. Two recent works illustrate this design philosophy in adjacent domains. Wang *et al.* introduced *MotorBeat* [2], a system that transmits status messages from small DC-motor-driven home appliances to nearby smart speakers by modulating the motor’s pulse-width-modulation (PWM) drive to produce electromagnetism-induced acoustic tones; the appliance requires no additional networking hardware. Sharma *et al.* extended this concept to airborne platforms with HUM [1], an acoustic side-channel for inter-drone communication that uses brushed coreless motors as transmitters under BFSK modulation,

and which addresses the formidable challenge of broadband propeller noise through biomimetic propeller designs and acoustic insulation.

The present work occupies a specific point in this design space. Rather than the brushed coreless motors of HUM or the small DC motors of MotorBeat, it uses the brushless DC (BLDC) motors that dominate quadrotor propulsion in the consumer, industrial, and research drone markets. BLDC motors offer two physical advantages: a large, mechanically resonant motor bell and strong rare-earth permanent magnets, both of which amplify the magnetostrictive acoustic emission produced when high-frequency content is injected into the drive waveform. The mechanism is closely related to the electromagnetic torque-ripple acoustic noise long studied in BLDC drive design [3], but here that emission is treated as a feature rather than a defect: it becomes the carrier of the data link.

This report is the first phase of a longer-term research effort whose ultimate objective is acoustic time-of-flight (ToF) ranging between cooperating drones, an extension of the cost-effective ToF localization techniques recently demonstrated for ground robots [6]. The drone-acoustic literature is growing rapidly along multiple axes [4], including passive detection of UAVs by their aerodynamic signature [5]; the present effort instead investigates active acoustic transmission as a deliberate communication channel. The contributions of this work are threefold: (i) an end-to-end transmitter and receiver implementation using only commodity hardware and a single added MEMS microphone; (ii) characterization of the BLDC motor’s acoustic emission across the 17–22 kHz band, including identification of a strong 19 kHz mechanical resonance; and (iii) demonstration of real-time BFSK packet decoding at distances representative of close-formation flight.

## 2 System Concept and Architecture

The communication mechanism rests on *magnetostriction*: the tendency of ferromagnetic materials to undergo small dimensional changes under a varying magnetic field. Injecting a high-frequency component into the BLDC drive waveform causes the motor’s ferromagnetic structure — stator iron, rotor back-iron, and motor bell — to vibrate at that frequency, with the bell acting as a passive acoustic radiator. An unmodified motor can therefore emit measurable ultrasonic sound when commanded appropriately. The end-to-end system pairs an ESP32-S3-driven transmitter, detailed in Section 2.1, with a single MEMS-microphone receiver, detailed in Section 2.2. The data link is unidirectional, and the receiver does not perform clock recovery against the transmitter — simplifications revisited in Section 4 as the principal limitations of the current design.

### 2.1 Transmitter Design

The transmitter is a single ESP32-S3 module driving a SimpleFOC Arduino Shield. The shield is used purely for its hardware capabilities — its three-phase MOSFET bridges and gate-driver circuitry — with no use of its onboard control intelligence or any field-oriented control (FOC) algorithm. The actuator itself is a DJI 2212-920KV BLDC motor, selected for its large-diameter motor bell and high-strength permanent magnets, both of which amplify the magnetostrictive emission; higher-impedance gimbal-class motors (e.g., the 2208 series) produced markedly weaker acoustic output. All control logic, including commutation and carrier generation, resides in custom firmware running on the ESP32-S3.

Motor commutation is implemented in open-loop sinusoidal mode. No rotor position encoder, Hall sensor, or back-EMF observer is used. This is a deliberate choice: the overwhelming majority of consumer and research drone platforms ship without rotor encoders, and any communication architecture predicated on encoder feedback would be impractical to retrofit. Slip (loss of synchronization between

the drive waveform and the rotor) was observed during testing and was corrected, when necessary, by raising the drive amplitude. Future iterations are expected to incorporate sensorless back-EMF tracking, now standard in commodity electronic speed controllers, but its absence in the present implementation is intentional rather than incidental.

The signal generation pipeline produces a composite three-phase waveform consisting of two superimposed components. The first is the conventional commutation waveform: three sinusoids of identical frequency and amplitude, mutually offset by  $2\pi/3$  rad, whose frequency sets motor rotational velocity and whose amplitude sets torque. The second is the data carrier: a second set of three sinusoids, again mutually offset by  $2\pi/3$  rad, but oscillating at the ultrasonic carrier frequency (nominally 19–22 kHz) and amplitude-keyed by the data signal. The two components are summed sample-by-sample in the same interrupt service routine (ISR) that updates the PWM duty cycles, yielding a per-phase waveform of the form

$$v_\phi(t) = A_c \sin(\omega_c t + \phi) + A_d(t) \sin(\omega_d t + \phi), \quad (1)$$

where  $\omega_c$  is the commutation angular frequency,  $\omega_d$  is the carrier,  $\phi \in \{0, 2\pi/3, 4\pi/3\}$  selects the phase, and  $A_c$  and  $A_d(t)$  are the commutation and data amplitudes, with  $A_d(t)$  FSK-modulated. The total amplitude was bounded at 90% of full scale and split 70%/30% between commutation and carrier in all experiments reported here.

The firmware runs on the ESP32-S3’s 160 MHz PLL clock, which is immune to dynamic frequency scaling that might otherwise introduce jitter. A FreeRTOS ramp task at 1 kHz updates the shared motor state. The composite waveform is computed inside an ISR triggered exactly on the PWM timer’s Timer-Equals-Zero (TEZ) event, using a pair of 256-entry sine lookup tables resident in DRAM (rather than in flash) to prevent flash-cache stalls during ISR execution. The PWM peripheral is configured in center-aligned UP/DOWN mode at an effective 80 kHz, which substantially reduces phase-current ripple compared to edge-aligned modulation. Output duty cycle is bounded between 5% and 95% to prevent saturation states that would distort the carrier.

## 2.2 Receiver Design

The receiver is built around a single SPH0641LU4H-1 ultrasonic MEMS microphone, selected for its frequency response well above 20 kHz, which natively outputs a 1-bit pulse-density-modulated (PDM) stream. The microphone is clocked at 4.608 MHz from the ESP32-S3’s 160 MHz PLL via the chip’s PDM-to-I2S peripheral, which applies a fixed  $64\times$  hardware decimation to produce a 16-bit, 72 kHz mono PCM stream on the internal I2S bus. Samples are delivered to firmware in 144-sample DMA chunks, one chunk every 2 ms.

Demodulation is performed by a Goertzel filter pair operating on each received chunk. The Goertzel algorithm, which computes a single-bin DFT recursively at  $O(N)$  cost, is well suited to a microcontroller deployment that needs power estimates at only two known frequencies. For a symbol duration of  $T_s$  at sample rate  $f_s$ , the per-symbol window is

$$N = f_s \cdot T_s, \quad (2)$$

with frequency resolution approximately  $1/T_s$ . At the operating point used throughout this work —  $T_s = 100$  ms,  $f_s = 72$  kHz — this gives  $N = 7200$  samples and a bin width of approximately 10 Hz. The decoder accumulates one symbol’s worth of samples, computes the Goertzel power at the mark and space frequencies, and decides the bit value by comparison: mark power exceeding space power

yields a 1, otherwise 0. The use of a power *ratio* rather than absolute power makes the decision robust to amplitude variation due to range, microphone orientation, and platform-internal acoustic loading.

Above the bit-decision layer, a structured packet protocol provides framing and error detection. Each frame consists of an 8-byte preamble of 0xAA (yielding 64 alternating mark/space symbols), a 16-bit sync word (0xD391), a 1-byte length field, a 1-byte sequence number, a payload of 0–127 bytes, and a 2-byte CRC-16/CCITT-FALSE checksum computed over the length, sequence, and payload fields. The preamble plays a critical timing role: because the receiver advances its Goertzel windows at a fixed cadence from the moment decoding begins, with no clock recovery against the transmitter, the alternating preamble offers 64 opportunities for the receiver’s window boundaries to drift into approximate alignment with the transmitter’s symbol boundaries before the sync word arrives.

The packet recovery state machine is implemented as a six-state pipeline: HUNT → HDR\_LEN → HDR\_SEQ → DATA → CRC\_HI → CRC\_LO. In the HUNT state, a 16-bit sliding shift register is compared against the sync word on every received bit; on match, the machine advances to HDR\_LEN. Subsequent states accumulate the requested payload bytes and the CRC; on CRC failure, the frame is discarded and the machine returns to HUNT. A per-frame timeout, computed from the announced length and current baud rate, guards against frames truncated mid-flight by transient noise. The receiver firmware additionally exposes a UART configuration console at 115200 baud that allows runtime adjustment of mark frequency, space frequency, symbol duration, and decoding mode — a feature that proved essential during empirical tuning, as discussed in Section 4.

### 3 Methodology and Results

Three experiments were conducted to characterize the system. In all cases, the transmitter consisted of a DJI 2212-920KV motor mounted on a stationary jig, driven by an ESP32-S3 + SimpleFOC stack as described in Section 2.1. For all experiments, the motor speed was set to 500 rad/s, which is a rough approximation of rotor speed at hover for the FlameWheel 550 drone the motor was sourced from. The receiver was an ESP32-S3 + SPH0641LU4H-1 microphone, with PDM output additionally bridged to a connected workstation via USB audio for spectrogram capture and visual verification.

#### 3.1 Experiment 1 — Baseline Distance Characterization

The motor was configured to transmit a continuous 10101010 pattern at 10 baud, producing an alternating sequence of mark and space tones, with no propeller attached. This isolated the purely magnetostrictive emission from any aerodynamic noise. Recordings were captured at five distances — 0, 1, 1.8, 2.8, and 5.6 m — in an acoustically uncontrolled lab environment with typical ambient HVAC noise. Spectrograms (Fig. 1) show clear, well-separated spectral lines at both mark and space frequencies through 2.8 m, with strong harmonic content visible above the carrier. At 5.6 m the ambient noise floor noticeably encroaches on carrier power, and harmonic structure begins to fade into the background, but the fundamentals at the mark and space bins remain visually distinguishable. The narrowband emission also reveals the motor’s tendency to radiate at multiple harmonic frequencies, a property that may be exploited in future multi-tone schemes.

#### 3.2 Experiment 2 — Acoustic Viability Under Propeller Load

A propeller was attached and the motor spun up to operating speed. To characterize the broadband propeller noise floor and identify the most favorable carrier frequencies, a sequence of single-frequency

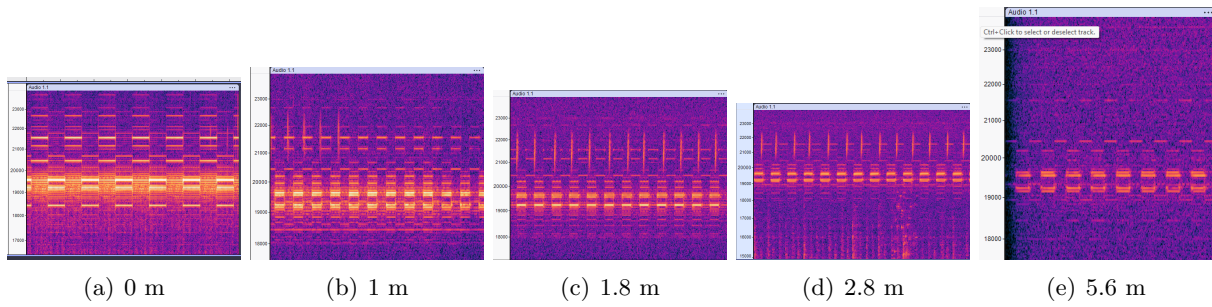


Figure 1: Baseline distance characterization: spectrograms of an alternating mark/space transmission with no propeller, captured at five separation distances between motor and microphone. Mark and space spectral lines remain distinguishable across the full range, with the noise floor encroaching most visibly at 5.6 m. The vertical axis spans roughly 17–23 kHz in each panel.

tone bursts was transmitted at 17, 18, 19, 20, and 21 kHz. The resulting spectrogram (Fig. 2) shows that for most of the swept frequencies, the tone is faint or imperceptible against the propeller’s broadband signature. The 19 kHz burst, however, induces a markedly bright and persistent acoustic feature, accompanied by visible harmonic structure, that stands out clearly from the noise floor — consistent with the motor exhibiting a strong mechanical resonance at this frequency. The exact physical origin of the resonance (motor bell, stator, propeller balance, or mounting) was not investigated in this phase.




Figure 2: Single-frequency tone burst sweep at 17, 18, 19, 20, and 21 kHz with the propeller attached and spinning at operating speed. Most frequencies are largely masked by broadband propeller noise; the 19 kHz burst (center panel) excites a strong mechanical resonance with visible harmonic structure that remains clearly distinguishable above the noise floor.

Motivated by this finding, a follow-up BFSK transmission was conducted with a deliberately wide mark/space split — 21 kHz mark, 17 kHz space — chosen to place the two tones on either side of the resonance peak, in order to maximize the visual contrast achievable under load. The resulting spectrogram is shown in Fig. 3. The transmitted bit pattern is identifiable by visual inspection of

the spectrogram, although the present Goertzel-only demodulator was not able to recover the data automatically with the present configuration. This result is interpreted as a strong indicator that on-air communication during powered flight is feasible in principle: the limiting factor is the receiver’s signal processing pipeline rather than the channel itself.




Figure 3: 10-baud BFSK transmission with the propeller running, using a deliberately wide mark/space split of 21 kHz / 17 kHz selected on the basis of the resonance characterization in Fig. 2. The alternating bit pattern is visible against the broadband noise floor; automatic recovery by the present demodulator was unsuccessful, suggesting the bottleneck lies in receiver-side signal processing rather than the acoustic channel.

### 3.3 Experiment 3 — End-to-End Packet Decoding

A complete packet — preamble, sync word, header, payload, and CRC — was transmitted in a quiet environment with the propeller again detached at 10 baud, with mark  $\approx 21.5$  kHz and space  $\approx 20.0$  kHz. Fig. 4 shows the captured spectrogram at 1.8 m from the source. The mark and space tones are bright, well-separated, and free of overlap with the harmonic skirt of the opposing tone, and the alternating bit structure of the transmitted pattern is plainly visible. Across multiple consecutive packet transmissions, the live receiver achieved *100% packet retention* after manual restart-based timing alignment, with all CRC-16 checks passing for the recovered payloads. This result establishes a baseline: a complete, bit-accurate, integrity-checked acoustic data link from a single drone-grade BLDC motor to a single MEMS microphone, at a separation representative of close-formation flight.

Aggregated, the three experiments characterize the channel as follows. The acoustic emission of an unmodified DJI 2212-920KV motor is detectable through approximately 5.6 m in still air. End-to-end packet recovery has been demonstrated through 1.8 m at 10 baud in stationary, no-propeller conditions. Under aerodynamic load, the motor exhibits a strong mechanical resonance at 19 kHz that remains spectrally distinguishable against propeller noise, and the observed spectrograms support the interpretation that a robust on-load link is achievable with an improved demodulator.

## 4 Discussion

Several artifacts of the present design merit discussion. First, a consistent frequency offset has been observed between the carrier as transmitted and the spectral peak as received, the latter typically displaced by tens to a few hundred Hz from the nominal carrier. The leading hypothesis is that the




Figure 4: Proof of concept: 10-baud BFSK transmission captured at 1.8 m, motor stationary, with mark and space tones at approximately 21.5 kHz and 20.0 kHz respectively. Both tones exhibit high spectral contrast against the noise floor, and the bit pattern is plainly visible. The live receiver decoded packets with 100% retention under these conditions.

spinning motor acts as a slowly time-varying acoustic filter — its mechanical resonance shifting with rotor angle, temperature, and load — but Doppler shift due to rotor tangential velocity, intermodulation between carrier and commutation, and residual receiver-clock error have not been excluded. In practice, this shift was compensated by manual retuning of the Goertzel filter bins via the runtime UART console; an autonomous deployment would require automatic peak-tracking similar to the adaptive bin-selection approaches employed in passive acoustic UAV detection [4].

Second, the receiver’s lack of clock recovery is the most serious systemic limitation of the current design. The 64-bit alternating preamble provides only the most rudimentary form of timing alignment; reliable initial lock typically required several manual restarts of the receiver firmware until its free-running symbol windows happened to align with the transmitter’s. A standard digital phase-locked loop on the bit transitions during the preamble would eliminate this manual step entirely and provide ongoing tracking against transmitter clock drift. Implementing such a PLL is the next engineering priority and is a prerequisite for the more sophisticated noise-rejection schemes [4] likely required for robust operation under powered flight.

Third, the present 10 baud raw rate is low; useful payload throughput is lower still after framing overhead. Increasing baud rate trades directly against the symbol-window length  $N$  in (2), which widens the Goertzel bin and demands greater frequency separation between mark and space — itself constrained by the motor’s narrow region of strongly-coupled emission and by the spectral purity of a 4-samples-per-cycle PWM representation. Forward error correction at the link layer is expected to recover much of the headroom currently consumed by retransmission and CRC failure, while faster switching hardware on the motor driver would unlock much more precise waveform construction for relatively little cost..

## 5 Concluding Remarks

This work has demonstrated that a standard BLDC drone motor can be repurposed as an ultrasonic acoustic data transmitter through the controlled exploitation of magnetostriction, with no addition of dedicated transmitter hardware to the platform. Bench measurements show clear acoustic emission through 5.6 m in stationary conditions, identify a high-SNR 19 kHz resonance that survives broadband propeller noise, and demonstrate live, real-time packet decoding with 100% retention at 1.8 m. The asymmetry of the system — a complex transmitter requiring no new hardware, paired with a receiver requiring only a single MEMS microphone — reduces the marginal cost of equipping an existing drone with this capability to essentially that of one microphone.

Several directions follow naturally. Closed-loop symbol timing recovery is the immediate priority. Beyond that, characterization of the observed frequency-offset phenomenon, exploration of On-Off Keying at the 19 kHz resonance as a candidate modulation more robust under load than BFSK, and the introduction of forward error correction at the link layer are all expected to extend operational range and reliability. The longer-term objective of the program, drone-to-drone time-of-flight ranging in the spirit of recent low-cost ToF localization for ground robots [6], will require improvements in timing precision and signal-to-noise ratio beyond those of the present architecture, but the channel demonstrated here is the foundation upon which that work can proceed.

## References

- [1] S. Sharma *et al.*, “Acoustic Side-Channel Communications for Aerial Drones with HUM,” in *Proceedings of the 22nd ACM Conference on Embedded Networked Sensor Systems (SenSys)*, 2024.
- [2] W. Wang *et al.*, “MotorBeat: Acoustic Communication for Home Appliances via Variable Pulse Width Modulation,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 1, 2022.
- [3] K. Xia *et al.*, “Acoustic Noise of Brushless DC Motors Induced by Electromagnetic Torque Ripple,” *Journal of Power Electronics*, vol. 17, no. 3, pp. 597–606, 2017.
- [4] C. Kang *et al.*, “From classical approaches to recent advancements: A holistic review of acoustic detection for unmanned aerial vehicles,” *AIP Advances*, vol. 15, 2025.
- [5] Z. Ghouli, “Passive acoustic detection and localization of drones using MEMS microphones and machine learning,” *Acta Acustica*, vol. 10, 2026.
- [6] B. M. Al-Hadithi and C. Pastor, “Cost-Effective Localization of Mobile Robots Using Ultrasound Beacons and Differential Time-of-Flight Measurement,” *Applied Sciences*, vol. 14, 2024.