Hark: Breaking Receiver Power Limits of LoRa

Amalinda Gamage, Pramuka Sooriya Patabandige, Spanddhana Sara and Ambuj Varshney
National University of Singapore
amalinda@nus.edu.sg, {pramuka, spanddhanasara}@u.nus.edu, ambujv@nus.edu.sg

Abstract

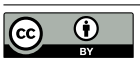
A long-standing challenge for LoRa devices has been to enable ultra-low-power reception that supports always-on listening while preserving the standard LoRa PHY and compatibility with existing COTS transmitters. Currently, COTS LoRa radios fail to support spontaneous actuation due to the prohibitive receive power (~43.5 mW) required to listen continuously, which effectively locks out applications that fall under LoRaWAN Class C. Recent works achieve substantially lower receive power, but either require modified physical layers and coordinated helper transmissions, or forgo dechirping and the range it provides. To this end, we present Hark, an ultra-low-power LoRa receiver that enables always-on listening at 1.90 mW while fully demodulating and decoding standard frames from unmodified COTS transmitters. We achieve this by pairing a novel digital wake-up pipeline with an analog tunnel-diode oscillator frontend that exhibits autodyning and is disciplined by a shared, unsynchronized, unmodulated carrier emitter. Specifically, Hark samples the channel at a fraction of the Nyquist rate to minimize compute and ADC power, yet successfully demodulates standard LoRa frames even below the noise floor. Hark demonstrates a communication range of nearly 500 m while still reducing receive power by over 22× compared to the COTS LoRa receiver. Hark therefore targets deployments where standard LoRa transmitters must remain unmodified and a receiver-side carrier emitter can be shared across multiple receivers.

CCS Concepts

• **Hardware → Wireless devices**.

Keywords

LoRa, ultra-low-power receiver, tunnel diodes.



This work is licensed under a Creative Commons Attribution 4.0 International License.
MobiCom '26, Austin, TX, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2505-0/2026/10
https://doi.org/10.1145/3795866.3844153

ACM Reference Format:
Amalinda Gamage, Pramuka Sooriya Patabandige, Spanddhana Sara and Ambuj Varshney. 2026. Hark: Breaking Receiver Power Limits of LoRa. In *The 32nd Annual International Conference on Mobile Computing and Networking (MobiCom '26), October 26–30, 2026, Austin, TX, USA*. ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3795866.3844153

1 INTRODUCTION

On-demand downlink communication for battery-powered devices is a critical feature that can push the boundaries of next-generation industrial and residential LPWANs. However, this feature is difficult to enable given the high energy cost of continuous listening that cannot be supported by battery-powered IoT devices.

Battery-powered IoT devices are now embedded in every corner of our world—in farms [64], factories [28], roads [5], rivers [6, 7, 9, 28], and homes [36]. They often operate in remote, infrastructure-poor environments where conventional short-range wireless networks fall short. To communicate across large, obstructed areas over the shared ISM spectrum, LPWAN technologies commonly trade off data rate to achieve long-range and low-power communication [16].

LoRa is widely recognized for its robust performance in challenging environments. For example, a LoRa radio can demodulate frames −20 dB below the radio noise floor [51]. This low-SNR capability, combined with other features such as concurrent transmissions on the same physical channel via orthogonal spreading factors and strong interference resilience, enables LoRa to deliver long-range communication at scale. Over the past decade, LoRa has evolved to become one of the most widely adopted LPWAN technologies, with over 125 million end-devices deployed globally [35]. When paired with aggressive duty cycling, a COTS LoRa transceiver can enable sensing applications to last years on a battery. Even under LoRa's most resilient configuration of SF-12, a single 16-byte uplink consumes under 250 mJ[1]. At one uplink per hour, total daily energy stays below 6 J, which a standard 0.7 Wh coin cell can sustain for over a year. As a result, LoRa has become the de facto modulation for wide area sensing, where devices only wake up briefly to transmit and stay asleep for the vast majority of time.

[1]The SX1262 draws approximately 38.5 mA when transmitting at +14 dBm under SF-12, measured using the Nordic Power Profiler Kit II [40].

Consequently, this offloads the burden of *continuous listening* for uplinks to LoRa gateways, which are typically grid-powered and not energy-constrained. This uplink-dominated architecture has proven to be remarkably energy efficient for sensing applications. But it breaks down when applications require immediate responses from devices that must listen continuously, not just speak.

Consider the following constraint: LoRaWAN operators often need to request urgent uplinks or push time-critical downlinks outside pre-programmed schedules. But most battery-powered devices operate under LoRaWAN Class A to conserve energy, which restricts reception to a short window immediately following an uplink. A downlink that arrives outside this window cannot be delivered because the radio is asleep. For example, actuating a smart water meter from a moving utility vehicle is impossible if the meter sleeps while the vehicle is in range. When the actuating gateway passes by, the transient opportunity to deliver a command or collect fresh data is lost. Even a simple status query suffers the same limitation. A user who asks "Is the valve open?" might wait hours for an answer under Class A. Such use cases are rare today not because they are unimportant, but because no commercial LoRa transceiver enables continuous listening without draining the energy budget.

The LoRaWAN standard defines Class C specifically to support such on-demand downlinks. However, the protocol is ahead of hardware capabilities. Today's COTS LoRa radios lack the efficiency to sustain the continuous listening that Class C requires. The SX1262, for example, consumes 43.5 mW in receive mode, which would demand large batteries or mains power. Battery-powered devices therefore operate within Class A in practice, where an unscheduled downlink must wait for the next uplink-triggered receive window.

Our design reduces this consumption to 1.90 mW. While still demanding on a coin cell, the reduction allows continuous listening for over 160 days on a pair of AA batteries, or indefinitely with a credit-card-sized solar panel. We envision applications in smart infrastructure as well as transient gateway scenarios, such as drive-by, walk-by, or fly-by interactions, where mobile gateways on cars, trains, drones, or satellites can trigger devices without prior scheduling. Ultra-low-power listening can therefore enable a new class of IoT that supports timely downlink delivery.

Low-power continuous listening can also augment Industrial IoT toward on-demand machine-to-machine control. These deployments already collect sensor data reliably but cannot act on it immediately without human intervention. A continuously listening end-device can change this. A gateway could query machine health sensors in a plant [26, 53], check fuel levels in storage tanks [24], trigger forest-fire alarms [33] on demand instead of waiting for the next uplink, or push an actuation command the moment an anomaly

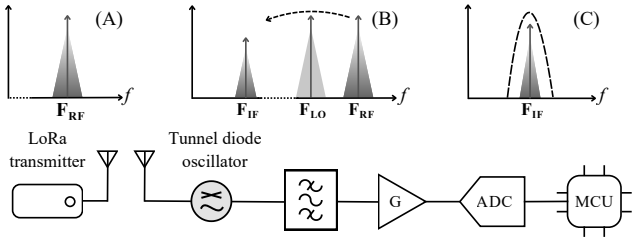


Figure 1: Overview of Hark: (A) LoRa frame arrives at F_{RF}. (B) The tunnel diode mixes it with its own carrier at F_{LO}, which an external emitter holds in injection lock, and the autodyning property of the diode produces an IF copy of F_{RF} at F_{IF}. (C) A bandpass filter isolates F_{IF}, which is then amplified, digitized, and demodulated.

is flagged, rather than dispatching technicians [41]. Reliable, low-latency downlinks can therefore turn passive sensing deployments into responsive ones that can ultimately help businesses scale and increase revenue [30].

Recent efforts, such as Sisyphus and μMote [54, 65], have demonstrated ultra-low-power reception by rethinking the receiver architecture. However, they make tradeoffs in compatibility or range to achieve such power budgets. μMote employs a custom frame structure where every data chirp must be accompanied by a reference chirp, which modifies the LoRa physical layer and breaks COTS compatibility. Sisyphus requires time-synchronized helper chirps from a secondary LoRa transmitter which significantly adds complexity, as each frame requires synchronized transmissions from two devices. While these prior works do minimize receiver power, they don't support direct compatibility with COTS devices.

Hark does not achieve the lowest receiver power among these efforts. If an application permits physical-layer modification, μMote reaches comparable range at 62 μW. But that range is delivered at one bit per symbol instead of SF bits per symbol, so the transmitter spends roughly an order of magnitude more airtime and energy per delivered byte. If per-frame synchronized helper transmissions are practical, Sisyphus offers 162.7 μW. But the helper occupies a second channel, which halves the spectrum available for data. Hark instead targets the case where receivers must accept standard LoRa frames from unmodified COTS transmitters without per-packet coordination. In exchange, Hark accepts a higher receiver power of 1.90 mW and a shared unmodulated emitter that needs no synchronization with the transmitter, and in return retains the full LoRa demodulation and decoding pipeline. This design is particularly relevant to deployments where existing LoRa transmitters and gateways should remain unchanged and shared receiver-side infrastructure can be amortized across many receivers.

Hark is a hardware–software co-designed LoRa receiver that (i) enables always-on listening at 1.90 mW, (ii) receives standard frames from unmodified COTS transmitters, and (iii) achieves nearly 500 m range. Hark contributes the following:

Microwatt Baseband Conversion. COTS LoRa receivers use power-hungry PLLs and mixers for downconversion [49]. Hark challenges this design by replacing these with an ultra-low-power autodyning frontend based on a tunnel diode oscillator (TDO) [37, 38]. While prior work has demonstrated autodyne receivers for narrowband reception such as FSK [37], it remains an open research question whether autodyning frontends can support complex modulations like LoRa, which require larger bandwidths and a digital demodulation pipeline that is key to demodulate frames buried below the noise floor. We experimentally demonstrate that an injection-locked autodyning frontend based on a single tunnel diode can successfully downconvert LoRa frames and track the demodulation performance of its COTS counterpart well into the negative SNR range, which has not been demonstrated before.

Sub-Nyquist Wake-Up Pipeline. We show that LoRa frame detection does not necessarily require full-rate sampling. We use this insight to design a digital wake-up pipeline that operates at a sub-Nyquist rate. We show that by exploiting a signature arising from the periodic nature of the LoRa preamble, our novel *multi-chirp stacking algorithm* recovers detection sensitivity even when signals are aliased and buried below the noise floor. This substantially reduces ADC power and MCU processing load, and decouples the higher energy cost of frame demodulation from frame detection.

Addressing. In dense deployments, receivers inevitably detect frames intended for nearby nodes. Invoking the full demodulation pipeline for every detected frame wastes significant energy. Hark minimizes this waste by discarding unintended traffic early. It achieves this using the NetID field, which is mandatory on all LoRa frames and supported by all LoRa radios. The two NetID symbols immediately follow the preamble and lack Gray-coding, whitening, or Forward Error Correction. Hark leverages this unencoded property to verify the address without invoking the heavy decoding pipeline. This yields a lightweight addressing mechanism that requires no modifications to the LoRa PHY.

Open-Source Demodulator. We implement a complete LoRa demodulator in C using the ARM CMSIS library for DSP operations. Unlike existing open-source demodulators such as gr-lora [45] and others [56, 69], which rely on compute-intensive DSP operations that are challenging to perform on low-power MCUs, our implementation is optimized for resource-constrained MCUs. In Section 3, we show that this demodulator matches the full SNR range of the COTS LoRa radio. We release this implementation to the community [17].

2 BACKGROUND

This section reviews the LoRaWAN link-layer protocol, the underlying LoRa physical layer, and the hardware constraints of commercial radios that motivate the design of Hark.

2.1 LoRaWAN Classes

The LoRaWAN protocol operates at Layer-2 of the OSI stack. It implements Media Access Control (MAC) and Logical Link Control (LLC) on top of the LoRa physical layer. LoRaWAN categorizes end-devices into three distinct classes based on the trade-off between power consumption and downlink latency [34]. Class A relies on uplink-triggered downlinks and is the most widely adopted. These devices spend the majority of their time in sleep mode and receive downlinks only during the short windows immediately following an uplink. This strategy minimizes power consumption at the cost of downlink latency, as the gateway cannot send data until the device initiates an uplink. Class B utilizes *beacon-based* downlinks, where devices open scheduled receive windows to receive periodic gateway beacons. It is a middle ground between Class A and Class C, where maximum downlink latency is set by the beacon period. Finally, Class C mandates *continuous listening* where devices remain in receive mode indefinitely. Consequently, Class C is currently feasible only for mains-powered devices; Hark aims to address this specific limitation for the battery-powered IoT.

2.2 LoRa Modulation

LoRa communication takes place over the shared ISM bands, such as the 868 MHz band. LoRa networks divide this band into multiple frequency channels, where each channel is utilized for both uplinks and downlinks.

A LoRa frame consists of multiple chirps that each sweep the channel bandwidth BW linearly. Each chirp spans one symbol duration of T_{sym}. LoRa uses three distinct chirps to build a frame. An *upchirp* ramps upward from $f_{\min}$ to $f_{\max}$, corresponding to a frequency sweep from $-BW/2$ to $+BW/2$ in baseband, where $f_{\max} - f_{\min} = BW$. A *downchirp* ramps in the opposite direction. A *data chirp* starts at an offset f_{sym} above $f_{\min}$, sweeps upward to $f_{\max}$, wraps back to $f_{\min}$, and continues until it completes a full sweep. The symbol duration is fixed across the frame and follows $T_{sym} = \frac{2^{SF}}{BW}$. As such, a higher spreading factor SF increases T_{sym} and thus the airtime. Increasing SF by one roughly doubles the energy to send the same payload. But in return, the LoRa frame becomes more robust against noise and interference, which significantly extends the communication distance.

A LoRa base upchirp, which linearly increases its frequency from $-BW/2$ to $+BW/2$ over one symbol duration, is given by:

$$C(t) = e^{j2\pi\left(f_0 t + \frac{1}{2}kt^2\right)}$$

where $f_0 = -\dfrac{BW}{2}$ is the initial frequency and $k = \dfrac{BW}{T_{sym}}$ is the gradient of the chirp. SF sets T_{sym} and therefore controls k. A modulated chirp is obtained by applying a frequency shift

f_{sym} to this base upchirp:

$$S(t; f_{\text{sym}}) = C(t) \cdot e^{j2\pi f_{\text{sym}} t} \tag{1}$$

The shift f_{sym} takes one of 2^{SF} values, so each chirp carries SF bits. To recover f_{sym}, the receiver multiplies each incoming chirp with the base downchirp $\overline{C(t)}$:

$$S(t; f_{\text{sym}}) \cdot \overline{C(t)} = e^{j2\pi f_{\text{sym}} t} \tag{2}$$

This process, known as *dechirping*, transforms the received chirp into a tone at f_{sym}. This achieves two things. First, when the receiver applies an N-point FFT to the dechirped samples, the index of the highest FFT peak maps directly to the originally transmitted f_{sym}. Second, dechirping concentrates the received symbol's energy, which was spread over the full channel bandwidth, into a single bin while noise remains distributed across all bins. This concentration is a key aspect that allows demodulating LoRa frames even when they fall below the noise floor.

A LoRa frame has four segments: a preamble of multiple base upchirps, two NetID chirps that encode the network ID, a Start Frame Delimiter (SFD) of $2\frac{1}{4}$ downchirps, and a sequence of data chirps that carry the payload. The preamble helps with packet detection and coarse synchronization. The NetID lets the receiver check whether an incoming frame is from an expected network. The end of the SFD marks the start of the first payload symbol and provides the sample-level timing accuracy needed for demodulation.

2.3 Bottleneck of COTS Receive Power

A COTS LoRa radio operates in three modes: transmit, receive, or sleep. We measure the power consumption of a Semtech SX1262 LoRa radio in each mode using the Nordic Power Profiler Kit II [40]. Fig. 3 shows the results. The radio draws 0.8 µA in sleep, 13.2 mA in receive mode, and 38.5 mA when transmitting a 16 B frame at 14 dBm under SF-12 at 3.3 V. Continuous listening draws about 1.6×10^4 times the sleep current. Over a 24 h day, continuous listening corresponds to 3.76 kJ ($\sim$ 1.05 W h), which would deplete a 0.7 W h standard coin cell in less than a day. The same coin cell can sustain hourly 16 B uplinks for over a year. This gap shows that while duty cycled sensing applications can last for over a year on tiny batteries, round the clock listening is not feasible.

To understand this gap, we look into the receive chain of a COTS LoRa radio in Fig. 2. In receive mode, the LNA continuously amplifies the incoming RF signal. A two port mixer, driven by a PLL based Local Oscillator (LO), downconverts the RF signal to baseband. For a LoRa frame within the 868 MHz band, the PLL tunes the LO to a nearby frequency in the same band so that mixing the two produces a low frequency baseband copy of the frame around 100 kHz. The baseband is then filtered, digitized by the ADC, and processed

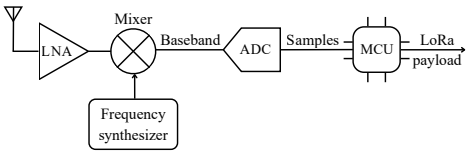


Figure 2: Simplified receiver chain of a COTS LoRa Radio.

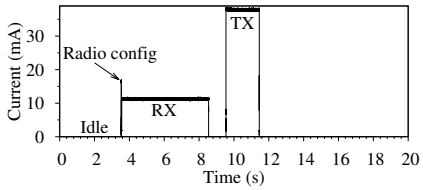


Figure 3: Current draw of SX1262 in sleep mode, in listening mode for five seconds, and transmitting 16 bytes at 14 dBm under SF-12.

by the MCU to perform detection followed by demodulation. Three blocks dominate receive power. From highest to lowest, these are the mixer and PLL used for downconversion, the LNA used for small signal amplification, and the baseband MCU that runs the LoRa PHY. We note that modern ADCs contribute very little to the total power, as 12-bit ADCs that can sample at four times the Nyquist rate needed for LoRa frames still operate at microwatt levels [58]. Significant reductions in overall receive power therefore require lowering the power consumption of the radio frontend and the baseband processor. But doing so is not straightforward. The LO and mixer greatly affect the purity of the baseband copy. Recall that LoRa demodulation depends on dechirping to concentrate the received symbol's energy into a single FFT bin, so the corresponding peak sits reasonably above the noise floor. If the LO is noisy or unstable, or if the mixer incurs high conversion loss, demodulation starts to fail at a much higher SNR which results in decreased range.

COTS LoRa radios do offer a Channel Activity Detection (CAD) mode, which confirms the presence (or absence) of an ongoing LoRa frame without invoking the full payload demodulation pipeline [16]. CAD briefly switches the radio into receive mode for several chirps and correlates the samples with a local chirp to determine whether a specific combination of channel and spreading factor is in use. While useful for CSMA, it is not a viable tool to enable continuous listening on battery-powered devices for two reasons: First, CAD draws 14 mW on average; invoking it repeatedly pushes system power into the tens of milliwatts, even before accounting for the MCU energy expended reconfiguring the radio for each CAD. Second, even if we resort to duty cycled CAD to save energy, for example performing CAD every couple of seconds, it still risks missing frames that start between those windows.

We therefore seek an alternative frontend that performs RF-to-baseband conversion at a much lower power budget

while preserving the spectral purity required for below-noise-floor LoRa demodulation. An injection-locked TDO exhibiting autodyning is one such candidate. We first review the TDO, its autodyning property, and prior work that uses it for long-range communication. We then quantify its down-conversion performance relative to a COTS frontend.

2.4 Tunnel Diodes as Microwatt Mixers

A tunnel diode is a semiconductor device with a heavily doped PN junction whose current–voltage curve exhibits a region of negative resistance. In a conventional diode, the forward current increases monotonically with forward-bias voltage. In a tunnel diode, the current first increases and then decreases as the bias voltage is increased further, as shown in Fig. 4. When biased within this negative resistance region, the diode behaves as a free-running oscillator. Coupling it with an LC resonant network lowers the frequency of oscillation and adds a degree of stability. The total power consumption of a TDO lies within the microwatt range.

Autodyning refers to the operating mode in which an RF oscillator simultaneously generates a local carrier and mixes that with incoming RF signals through its nonlinearity to produce Intermediate-Frequency (IF) outputs. TDOs by default operate in the autodyning mode. Therefore, employing a single TDO instead of two separate devices for LO generation and for mixing can help bring down the total power of an RF frontend from milliwatts to microwatts. This reduction in power is not without cost, as TDOs trade off stability for low power consumption. As a result, a free running TDO has higher phase noise and frequency drift as compared to off-the-shelf PLL based LO generators [62]. Additionally, TDO stability can also be affected by environmental factors, e.g., fluctuations in temperature, humidity [60].

Most oscillators are inherently subject to injection locking. This allows them to couple and synchronize with each other [10, 39]. TDOs too can therefore be stabilized this way. Prior work shows that a precision injection-locking carrier significantly improves the phase noise and frequency stability of a TDO [63], which is consistent with classical injection-locked oscillator behavior reported for other RF oscillators [8]. In this configuration, a separate emitter transmits the carrier, and the TDO locks to it even when the emitter is located at a considerable distance from the receiver. This is because the strength of the injection-locking carrier does not strictly determine the power of the LO. The TDO is an active oscillator that produces its own tone; the carrier merely disciplines its frequency and phase. Unlike helper signals which may require per-frame synchronization and customization, this carrier is a static, unsynchronized single tone which can discipline multiple TDOs across a large deployment area.

These properties have motivated the use of tunnel diodes across a large body of RF works. Prior works have employed tunnel diodes as a key component to build transmitters that achieved communication distances of several kilometers while operating at microwatt power [1–3, 11, 15, 27, 60–62]. SoMix [37] is the first to demonstrate low-power reception with an autodyning tunnel diode frontend using 2-FSK. Under LoS experiments, it receives 2-FSK frames transmitted at 30 dBm over 125 m while the frontend consumes only 100 μW. For its 2-FSK operation, an incident emitter strength of −70 dBm at the receiver was sufficient to hold the TDO injection locked. Hark builds on this frontend to enable COTS-compatible LoRa reception.

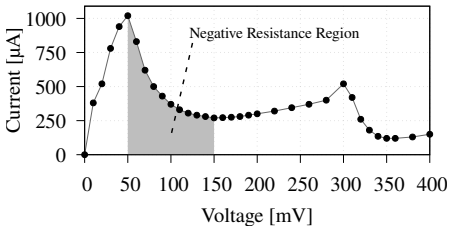


Figure 4: Current–voltage characteristics of the *GE1N3712* tunnel diode. The shaded negative-resistance region, where current decreases with increasing voltage, is the operating region of TDOs.

Cost and Availability. While tunnel diodes currently cost more than conventional diodes, this is a byproduct of low production volumes rather than fabrication complexity. Because their internal architecture is intrinsically simple, they were successfully mass-produced at scale decades ago [12]. They were phased out because their extremely low-power nature was a disadvantage at the time. That same characteristic now makes them highly desirable for ultra-low-power transceivers. As recent RF research continues to demonstrate their value for microwatt reception [1, 2, 37, 60, 62, 63], the renewed demand could motivate large-scale manufacturing.

3 ENABLING LOW-POWER BASEBAND

Prior works have demonstrated that injection-locked TDOs can serve as low-power frontends for simple modulation schemes such as OOK and FSK. However, LoRa modulation is more demanding, and successful demodulation depends on the spectral purity of the final baseband signal. We must therefore answer a fundamental question: *Could a TDO frontend preserve the baseband spectral fidelity needed for LoRa demodulation?* Our results show that it tracks the receive performance of a COTS radio with negligible reduction in SNR.

We establish this with a cabled experiment in which the TDO frontend and the reference COTS receiver observe identical signal and noise conditions across all iterations. As shown in Fig. 5b, the SX1262 transmitter sends standard LoRa frames on SF-12 at F_{RF}. The USRP B200 generates two signals on the same output: a CW tone that serves as the injection locking carrier, and band-limited Gaussian noise across the LoRa channel. Both the LoRa transmitter and the USRP feed

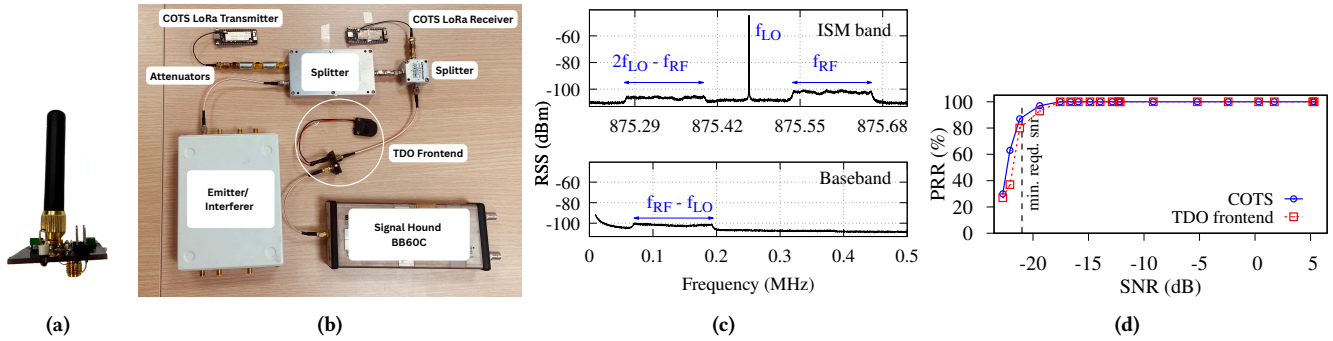


Figure 5: Demodulation performance of the TDO frontend: (a) The TDO frontend. (b) Cabled setup to deliver identically attenuated LoRa frames to both the TDO frontend and the COTS SX1262 receiver. (c) Frequency downconversion: the incoming LoRa frame at F_{RF} = 875.6 MHz mixes with the TDO carrier at F_{LO} = 875.468 MHz within the TDO (top), and produces a low-frequency IF copy centered at 132 kHz (bottom) that is digitized for baseband demodulation. (d) PRR under varying SNRs shows that the TDO frontend can track the COTS performance well into negative SNRs.

into an RF combiner. The combined signal passes through attenuators and into an RF splitter whose outputs feed both the reference SX1262 COTS LoRa receiver and the TDO frontend.

The IF output of the TDO frontend was fed to a Signal Hound BB60C configured in SDR mode. The BB60C was tuned to a center frequency of 150 kHz with a sampling rate of 650 kSPS to digitize the IF output and produce IQ data. Fig. 5c visualizes this downconversion. The LoRa frame at F_{RF} and the TDO's carrier at F_{LO} are both visible in the RF spectrum alongside a second-order mixing product. Below the RF band, a clean IF copy of the frame appears at $F_{RF} - F_{LO}$. This is the signal Hark filters, amplifies, and demodulates.

We sweep through LoRa's effective SNR range by adjusting the USRP noise power, the transmitter output power, the attenuation, or a combination of the three. At each SNR point, the transmitter sends 10 frames. We repeat this procedure for 10 runs. This yields 100 frames per SNR. The SX1262 receiver reports the SNR for each received frame under the same conditions. We use the averaged reported SNR as ground truth for both paths. We compute PRR for the TDO path by processing the recorded IQ traces with our custom C-based LoRa demodulator, the same codebase that runs on the MCU in the full Hark prototype. We compute the PRR of the COTS path directly from the SX1262 reception logs.

Fig. 5d plots PRR against SNR for both the COTS receiver and the TDO based pipeline. The results show that the tunnel diode frontend tracks the COTS performance closely. Both curves exhibit the waterfall drop-off at similar SNR values. This indicates that autodyning introduces a negligible SNR penalty and well preserves the signal integrity required for LoRa demodulation. The frontend is therefore unlikely to be the dominant sensitivity bottleneck in our prototype. Constraints will likely come from subsequent analog stages or low-power signal processing requirements. This validation motivates the specific design choices of Hark described next.

4 METHODOLOGY

Hark targets continuous listening and aims to achieve three design goals. First, it must be COTS compatible. Second, idle listening power must be low enough to run the receiver for months on a small battery or indefinitely from a credit-card-sized solar panel. Third, it must retain the key features of LoRa demodulation such as below noise floor operation.

Benefits of COTS Compatibility. LoRaWAN already benefits from extensive global adoption, which includes over 125 million deployed LoRa nodes [35] connected to both terrestrial and non-terrestrial gateways. Therefore, being fully COTS compatible allows for faster integration with this ecosystem. In contrast, being merely *LoRa-like* falls short of full compatibility, which slows adoption. We define COTS compatibility as receiving a standard LoRa frame from a transmitter that is oblivious to requirements outside the LoRa PHY.

4.1 Design of the Frontend

Hark's frontend converts an incoming LoRa frame at F_{RF} to a low-frequency IF signal suitable for digital processing. Fig. 6 shows the complete signal chain. The chain consists of six stages: baseband conversion, bandpass filtering, amplification, coupling, sampling, and baseband processing.

The tunnel diode in the TDO frontend is biased within its negative resistance region and coupled to an LC resonant circuit tuned within the sub-1 GHz band. An injection-locking carrier disciplines the TDO as described in Section 2.4, and this way the TDO generates a stable local carrier at F_{LO}. When an incoming LoRa frame arrives at F_{RF}, it is mixed with the TDO's F_{LO}. Hark focuses on the IF copy of the LoRa frame centered at $F_{IF} = F_{RF} - F_{LO}$, which we tune to fall at approximately 100 kHz by making a slight adjustment to F_{RF} and F_{LO}. Because the TDO's output contains the desired F_{IF} alongside harmonics and intermodulation products, we include a multiple-feedback bandpass filter tuned to 100 kHz with a 130 kHz passband to suit LoRa's 125 kHz operation.

This ensures that spectral content outside F_{IF} is highly attenuated before reaching the ADC. Without this filter, these components would alias into F_{IF} and degrade both detection and demodulation.

This filtered IF signal is weak. The TDO frontend incurs 25–30 dB of conversion loss [37], and the LoRa frames Hark targets operate well below the noise floor. Three low-power opamp stages amplify the IF signal before digitization. We set the gain so that the amplified noise floor alone swings across the full input range of 0 to 1.8 V of the ADS7042 ADC, covering all 12 bits even when no frame is present. The rationale behind this approach is that, at negative SNRs, the LoRa signal is buried in noise and is indistinguishable from it in the time domain. The signal would only reveal itself after dechirping concentrates energy into a single FFT bin. Fixing the system gain to serve strong positive SNR signals would mean significant degradation for frames below the noise floor. Setting the gain to fill the ADC range at rest ensures that weak LoRa frames are digitized using the maximum ADC resolution.

The amplified signal is then AC-coupled to remove the DC offset introduced by the amplifier. A resistive divider recenters the signal at the midpoint of the ADC range (0.9 V). This ensures that the bipolar IF signal is symmetrically placed within the unipolar ADC range. The external 12-bit ADC digitizes this IF signal. The MCU drives the ADC over SPI and collects samples at the rate required by the operating mode: sub-Nyquist rate during continuous detection, full Nyquist rate or higher during demodulation. Because sampling modestly above Nyquist has been shown to significantly improve demodulation performance [29, 69], Hark samples at 2× Nyquist during demodulation. The MCU receives the digitized samples and performs detection, synchronization, and demodulation. All DSP operations strictly use the ARM CMSIS library.

IF Gain Control. Hark omits amplification at RF. The antenna feeds the tunnel diode directly, so all amplification happens at IF. Recall that we set this gain against the noise floor so that all 12 ADC bits are utilized. At higher SNRs the same setting would drive the signal beyond 1.8 V, so the gain is lowered to keep the IF signal within the ADC range. We set this manually during experiments. Automating this task is left for future work.

Saturation at High SNRs. Under high received signal strengths, two distinct saturation limitations can occur. First, for a given IF-gain setting, a frame arriving well above the expected signal level can drive the amplified IF beyond the ADC input range, causing clipping and demodulation failure. Second, at much higher input power, the incident RF swing can drive the tunnel diode's instantaneous operating point outside the negative-resistance region in Fig. 4, degrading injection locking and autodyning. At still higher incident power, the RF excitation can drive the tunnel-diode junction current beyond its absolute maximum rating. Because tunnel diodes have extremely small junction areas, such overcurrent can produce rapid localized heating and permanent device damage [47].

4.2 Enabling Continuous Listening

To decouple frame detection from full demodulation, we build a detector that listens continuously at a sub-Nyquist rate and triggers full-rate sampling only once a frame is confirmed. The key insight is that Nyquist-rate sampling is necessary only for symbol recovery, not for frame detection. We show that undersampling by a factor of 4 still detects frames reliably, even those below the noise floor.

Undersampling reduces both ADC and MCU power consumption as lower rates cut ADC current and proportionally reduce the MCU processing load. For example, the ADC used in our system consumes 272 µW at 1 MSPS. Sampling at a fraction of that rate would reduce ADC power draw to tens of microwatts. If detection remains reliable at such low rates, continuous listening would become feasible within extremely small energy budgets.

However, sub-Nyquist sampling causes aliasing, which severely degrades demodulation or makes it infeasible. This does not affect our approach, since our goal is to detect frames at the lowest energy budget and trigger full-rate sampling only once a frame is confirmed. To make undersampled detection practical, two challenges need to be addressed. First, we must work with aliasing-induced ambiguities and still achieve reliable detection. A false negative means a missed frame. A false positive wastes energy and triggers unnecessary MCU processing. This challenge becomes further complicated by radio impairments such as CFOs and SFOs. Second, even when a frame is detected, it may not be addressed to that device, and proceeding to demodulate such frames wastes energy. In what follows, we present a methodology that addresses these challenges.

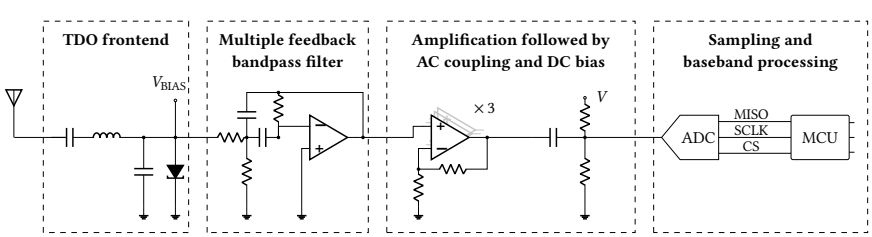
Undersampled Chirps. Consider a single transmitted upchirp $S(t; f_{\text{sym}})$. When $S(t)$ travels through a wireless channel, the received signal can be expressed as:

$$r(t) = h(t)\, S(t; f_{\text{sym}})\, e^{j2\pi\Delta f_{\text{cfo}} t} + n(t),$$

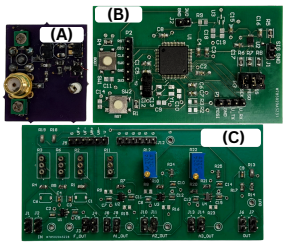
where $h(t)$ denotes the channel-induced amplitude and phase variations, Δf_{cfo} is the CFO, and $n(t)$ is noise. A typical LoRa receiver would sample this chirp at Nyquist-rate by setting $F_s = BW$. And the discrete-time representation of the received chirp can then be written as:

$$R(k) = r\left(\frac{k}{BW}\right), \quad k = 0, 1, \ldots, 2^{SF} - 1.$$

However, if the receiver reduces the sampling rate by a factor of D, the effective sampling frequency becomes $F_s = BW/D$. Now, the corresponding discrete-time sample



(a) Hark receiver chain



(b) Hark PCB prototype

Figure 6: The incoming frame is downconverted by the TDO frontend (A), bandpass filtered at IF with a multiple-feedback filter, amplified by three identical stages, AC-coupled, DC-biased, digitized, and fed to the MCU over SPI (B,C).

representation of the undersampled chirp can be written as:

$$R_D(k) = r\left(\frac{Dk}{BW}\right), \quad k = 0, 1, \ldots, \frac{2^{SF}}{D} - 1. \tag{3}$$

Thus, undersampling by D reduces samples per chirp by D, even though the chirp still sweeps the full $[-BW/2, BW/2]$.

According to the Nyquist–Shannon sampling theorem, a signal sampled at F_s can only be reconstructed in the range $[-F_s/2, F_s/2]$. Should a LoRa receiver sample this chirp at the sub-Nyquist rate $F_s = BW/D$, then the original chirp, which spans the full $[-BW/2, BW/2]$, will be folded into $[-BW/(2D), BW/(2D)]$. We show this aliasing effect using a single LoRa upchirp in Fig. 7. The figure compares an SF-12 upchirp of 125 kHz bandwidth at the Nyquist rate with undersampled versions at factors of 2 and 4. We emulate undersampling by a factor of D by retaining one out of every D Nyquist samples. At the Nyquist rate (Fig. 7a), the chirp appears as a clean monotonic frequency sweep across $[-BW/2, BW/2]$. At lower rates (Figs. 7b–c), the sweep remains monotonic, but components above $BW/(2D)$ fold back into the reduced band. This folding creates the appearance of several segmented chirps sweeping the compressed BW/D range within one symbol duration. The chirp structure is retained but compressed into a smaller frequency span. Our detector leverages this compressed representation to confirm the presence of a LoRa frame at far lower power cost than at Nyquist-rate listening.

Dechirping Undersampled Chirps. Dechirping is the key step that enables a LoRa receiver to confirm the presence of a frame. Existing LoRa implementations that operate at the Nyquist rate [45, 56, 69] rely on continuous dechirping to detect the preamble. Recall that a preamble consists of multiple repeated upchirps (e.g., 8–12 upchirps). As the receiver dechirps and FFTs these symbols, the dominant bin appears at the same index across consecutive preamble chirps. Tracking this bin is a widely used technique to confirm the presence of a LoRa frame. Although bin magnitudes vary slightly due to noise, the FFT bin index remains consistent. This consistency provides strong statistical evidence that a LoRa frame is present, even below the noise floor [16, 31]. However, this strategy does not extend cleanly to undersampled receivers, where aliasing reduces SNR and makes consecutive bin alignment unreliable.

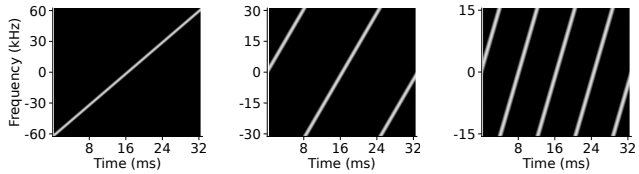
Impact of Undersampled Dechirping. Upon dechirping, the received LoRa symbol reduces to a single sinusoid plus noise [67]. If a symbol has N samples, the dechirped sinusoid accumulates in a single FFT bin. The signal amplitude in that bin therefore grows in proportion to N. Since noise is uncorrelated and spreads across all bins, its rms amplitude in any given bin grows only as $\sqrt{N}$ [59]. As a result, the FFT signal power scales as N^2 while the noise power scales as N, making the post-FFT SNR proportional to N. With undersampling by D, only N/D samples remain, so post-FFT SNR drops by D and the peak weakens proportionally. This makes reliable detection harder at lower sampling rates and motivates our multi-chirp stacking discussed next.

Multi-Chirp Stacking. To restore detection reliability under undersampling, we stack multiple chirps from the preamble. The receiver continuously collects an M-symbol window of undersampled samples while listening. Because preamble upchirps are identical and repeat every symbol duration, this window captures the same waveform M times even without prior symbol synchronization. It processes these M chirps in two groups of $M/2$ chirps. For the first group, it sums the k-th samples of chirps 1 through $M/2$ to produce one composite sample. Repeating across all samples yields one stacked chirp synthesized from $M/2$ undersampled chirps. The same process applied to chirps $M/2 + 1$ through M yields a second stacked chirp.

Formally, let $N' = 2^{SF}/D$ be the samples per symbol when undersampling by D. Let $R_D^{(m)}[k]$ denote the k-th sample of the m-th undersampled chirp, with $k = 0, \ldots, N' - 1$ and $m = 1, \ldots, M$. We form two stacked chirps by summing corresponding samples:

$$\tilde{R}_1[k] = \sum_{m=1}^{M/2} R_D^{(m)}[k], \qquad \tilde{R}_2[k] = \sum_{m=M/2+1}^{M} R_D^{(m)}[k].$$

Thus, from M chirps we obtain two stacked chirps, each with higher effective SNR due to coherent addition. Coherent addition across $L = M/2$ chirps increases the chirp amplitude



(a) (b) (c)

Figure 7: Frequency-domain representation of a LoRa upchirp at the Nyquist rate (a) and undersampled by factors of 2 (b) and 4 (c).

by L, while uncorrelated noise grows only by $\sqrt{L}$. Consequently, the FFT-peak SNR of each stacked chirp improves by a factor of L:

$$\mathrm{SNR}_{\mathrm{stack}} \approx L \cdot \mathrm{SNR}_{\mathrm{single}}$$

We then dechirp each stacked chirp using a locally generated undersampled downchirp and take its FFT. Let $C_D[k]$ denote the undersampled downchirp. We compute

$$y_i[k] = \tilde{R}_i[k]\, C_D[k], \qquad Y_i[b] = \mathrm{FFT}_{N'}\{y_i[k]\}, \quad i \in \{1, 2\}.$$

Let $b_i^\star = \arg\max_b |Y_i[b]|$ denote the peak bin of the i-th stacked chirp. We declare a frame if the two peaks lie within a small bin tolerance:

$$|b_1^\star - b_2^\star| \leq \tau.$$

This approach provides two advantages. First, stacking suppresses noise since uncorrelated noise averages out while the chirp adds coherently. Second, the agreement check between two stacked chirps reduces false positives, which is critical under undersampling where aliasing can mimic valid peaks.

Addressing. In dense deployments, nodes inevitably overhear frames intended for others. Fully demodulating such a frame only to discard it later wastes energy. So Hark uses the *NetID* field to reject unintended frames before full demodulation. The NetID is mandatory on all standard LoRa frames and is carried in two symbols that immediately follow the preamble. Unlike payload symbols, NetIDs are not Gray-coded, whitened, or protected by Forward Error Correction. This absence of encoding lets us recover an address without invoking the full demodulation and decoding pipeline. The result is a lightweight addressing mechanism that needs no modification to the LoRa PHY.

Together, the two NetID symbols directly represent a 16-bit sync word configured at transmission. However, without the redundancy of coding, this address space is difficult to use reliably. We therefore restrict the effective address space to a smaller subset. We can then accept both NetID peaks within a small tolerance of their expected bins, which makes address matching more resilient to interference.

Once a frame is detected, the receiver performs time synchronization and CFO correction as usual. It then demodulates only the two NetID symbols. If the recovered NetID matches the configured address within the tolerance, the receiver proceeds with full payload demodulation. Otherwise, the frame is discarded immediately.

5 PRIOR WORK

This section reviews ultra-low-power LoRa receivers that closely relate to Hark. We conclude with a summary of their performances in Table 1.

Low-Power LoRa Radios. Prior work has explored ultra-low-power LoRa receiver designs by rethinking the three fundamental steps in a LoRa receiver pipeline: (1) RF-to-baseband downconversion, (2) frame detection, and (3) demodulation. Each work proposes a novel approach to reducing the power consumption of one or more of these steps. But reaching these power budgets costs range, COTS compatibility, or both. What differentiates these designs is how much of each they trade off.

Saiyan [18] is a LoRa receiver that avoids the conventional dechirping process entirely. It transforms frequency-modulated LoRa chirps into amplitude-modulated signals using a surface acoustic wave (SAW) filter, followed by an envelope detector to bring the signal to baseband. Subsequent processing operates directly on this envelope, using peak detection to recover the data embedded in the original chirps. Saiyan requires no modifications to the transmitted signal, which preserves COTS compatibility. However, Saiyan bypasses dechirping, eliminating the processing gain that gives LoRa its long range, the very gain that motivates CSS modulation in the first place. Saiyan is therefore susceptible to interference, and its range is limited to around 55 m at 93.2 µW.

μMote [54] is the first work that exploits passive dechirping to build a LoRa-like receiver. It achieves a range of 400 m while consuming just 62 µW. However, μMote expects a modified LoRa frame with two chirps per symbol: one carrying data and the other serving as a reference for dechirping. This dual-chirp approach is clever in that it shifts the energy cost of generating the downchirp to the transmitter. But consequently, it relies on transmitter modifications that standard LoRa radios do not support. Moreover, μMote's modified frame encodes only a single bit per symbol rather than SF bits per symbol encoded in the LoRa PHY. Consequently, the transmitter also spends far more airtime and energy per delivered byte.

Sisyphus [65] is a passive receiver that achieves a line-of-sight range of 73 m at a continuous power budget of 162.7 µW, extendable to 107 m with an optional IF amplifier. However, Sisyphus shares a key limitation with μMote: both require modifications to the transmission scheme to aid reception. Sisyphus achieves pseudo-compatibility by using two synchronized commercial transmitters—one sending the payload and the other transmitting helper chirps. This approach introduces practical deployment challenges. First, receiving a single frame requires two separate COTS devices operating in precise time synchronization. Second, it breaks LoRaWAN

Table 1: Comparison of ultra-low-power LoRa receivers and backscatter systems. Sensitivity in dBm.

Work	Sensitivity(dBm)	Range	Power	PHY mod.?
Low-Power Receivers (Downlink)				
Hark (Ours)	−62‡	∼500 m	1.90 mW	No
Saiyan [18]	−85.8	∼55 m	93.2 µW†	No
μMote [54]	−48	400 m	62 µW	Yes (TX)
Sisyphus [65]	−36	73–107 m	162.7 µW	Yes (TX)
Backscatter Transmitters (Uplink)				
LoRa Backscatter [55]	–	0.5–2.8 km	9.25 µW	No
PLoRa [42]	–	1.1 km	2.59 mW	Yes (RX)
ALoBa [19]	–	50–200 m	0.3 mW	Yes (RX)
NetScatter [20]	–	Office	µW level	Yes (RX)
P^2LoRa [66]	–	2.2 km	320 µW	Yes (RX)

†A simulated figure. ‡95% PRR at $n/4$.

compatibility, as the receiver can only decode frames sent in this coordinated manner.

Delegating carrier generation to an external source is common practice among low-power architectures. Compared to backscatter and Schottky-diode receivers [46, 55], Hark only requires the incident emitter signal to stay above −57.3 dBm, and its carrier requires no synchronization with the transmitter, carries no data, and a single emitter serves a large deployment of multiple receivers. Unlike Sisyphus, Hark receives from any unmodified COTS LoRa transmitter without coordination. Finally, the helper signal in Sisyphus occupies the adjacent channel, which reduces spectral efficiency.

Downclocked Receivers. Prior undersampled receivers include E-MiLi [72], which reduces WiFi idle-listening power via sub-Nyquist detection but requires a custom preamble on every frame which calls for driver-level changes at the transmitter. LiteNap [68] demodulates LoRa at sub-Nyquist by resolving aliasing via hardware phase jitter fingerprints. While Hark shares with both the core premise that reception need not run at a single rate, it differs in two major ways. First, both utilize I/Q data from an SDR frontend which limits the lifetime of battery-powered applications due to high power. Hark instead relies on far lower fidelity baseband from the ultra-low-power tunnel diode frontend, yet reaches a longer range than LiteNap and demodulates frames at negative SNRs, which LiteNap does not report. Second, prior works process samples on a host PC, whereas Hark is a hardware-software co-design whose processing fits within the compute and energy budget of an MCU. These constraints drive the design of Hark and are non-negotiables for the class of devices it aims to serve.

Wake-up Radios. These reduce idle power by pairing the main radio with an ultra-low-power receiver that wakes it only when a transmission is detected. To maintain microwatt energy budgets, the second low-power receiver is designed to rely on simpler modulations like OOK or FSK. Works like eWake [25] and Piyare et al. [43] exploit the multi-modulation capability of LoRa transceivers by reconfiguring COTS radios to transmit OOK beacons to a secondary wake-up receiver. But the newer LoRa radios, e.g., the SX1262 and LR1121, no longer offer OOK [51, 52]. Even where supported, this requires the transmitter to signal wake-up calls outside the LoRaWAN standard which contrasts with Hark's compatibility goal.

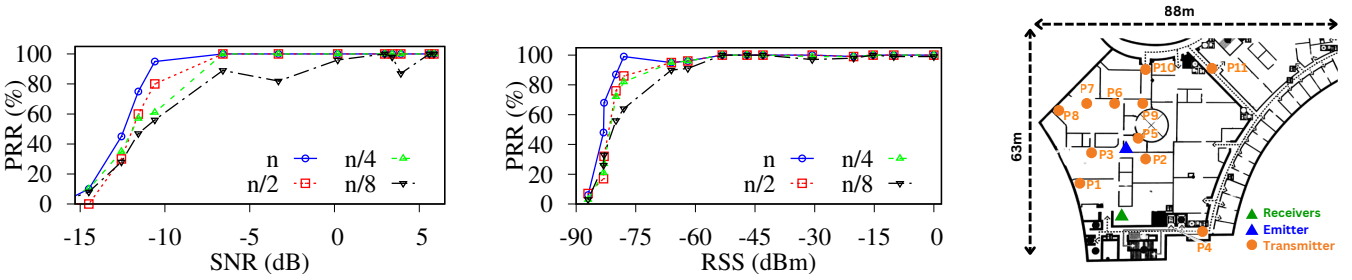
6 EVALUATION

We evaluate Hark through cabled, indoor, and outdoor experiments and conclude with a detailed power breakdown. Cabled experiments isolate Hark from interference and give precise control over SNR, while the indoor and outdoor experiments test it under real deployment conditions.

6.1 Cabled Experiments

We reuse the cabled setup of Section 3 and replace the Signal Hound BB60C with the full Hark prototype. Everything else remains unchanged. We sweep SNR and received power as before. The COTS receiver reports the SNR and RSSI of each frame, which we use as ground truth. To evaluate Hark over hundreds of frames per SNR point, we tap the SPI bus between the Hark MCU and the ADC using a Saleae logic analyzer. The MCU remains in full control throughout. It continuously samples at the sub-Nyquist detection rate, and upon detecting a frame, autonomously ramps the ADC to full demodulation rate. The logic analyzer passively records every sample crossing the SPI bus. We process the recorded traces offline using the same C codebase that executes on the MCU, and verify identical demodulation outcomes between the two methods. We follow this procedure across all experiments in this section.

Fig. 8a plots PRR against SNR for four sampling rates: Nyquist (n), $n/2$, $n/4$, and $n/8$. Rates n through $n/4$ achieve 100% PRR down to approximately −6.5 dB SNR. Below this point, the curves diverge. The Nyquist-rate detector maintains 95% PRR at −10.6 dB, whereas the $n/4$ detector drops to roughly 60% at the same SNR. At $n/8$, performance degrades earlier and less predictably. Based on these results, we select $n/4$ as Hark's default detection rate. This reduces ADC power and MCU processing load by a factor of four while maintaining reliable detection well into LoRa's below noise floor operation. Reducing further to $n/8$ offers diminishing returns, with marginal power savings and significantly worse detection reliability. Fig. 8b plots PRR against RSSI. We sweep up to 0 dBm, the SX1262's maximum recommended input power [51], and observe that Hark maintains approximately 100% PRR at high RSS values, which implies that Hark's frontend does not saturate within the intended operating regime. Hark achieves a sensitivity of −62 dBm at $n/4$, which we define as the minimum RSSI required to achieve 95% PRR. PRR degrades gracefully below this point and remains above 70% down to −80 dBm.



(a) PRR vs. SNR (b) PRR vs. RSSI (c) Indoor floor plan with positions

Figure 8: (a, b) Cabled evaluation of Hark at four sampling rates (n, $n/2$, $n/4$, $n/8$), up to SX1262's maximum limits. (a) Rates n through $n/4$ hold 100% PRR down to −6.5 dB SNR. (b) At $n/4$, Hark reaches 95% PRR at −62 dBm and 70% at −80 dBm. (c) Indoor floor plan with positions P1–P11.

Table 2: Indoor PRR (%) of Hark across four sampling rates. Shaded rows show the effect of increased TX power for identical positions.

Position	RSSI	n	$n/2$	$n/4$	$n/8$
TX power: +22 dBm					
P1	−40.8	100	100	100	99
P2	−55.6	100	97	95	91
P3	−57.7	100	100	98	96
P4	−65.0	94	92	90	84
P5	−67.8	93	89	81	76
P6	−69.6	95	90	78	52
P7	−73.6	91	87	72	49
P8	−67.9	83	76	65	48
P9	−75.7	0	0	0	0
P10	−82.9	0	0	0	0
P11	−88.6	0	0	0	0
TX power: +30 dBm					
P8	−52.7	100	100	99	94
P9	−56.3	100	98	95	90
P10	−62.8	67	43	26	15
P11	−64.7	96	93	87	72

6.2 Indoor and Outdoor Range

Indoor. We evaluate Hark in an indoor NLoS environment within a university lab which spans multiple rooms separated by concrete walls, doors, and corridors. Fig. 8c shows the floor plan and positions.

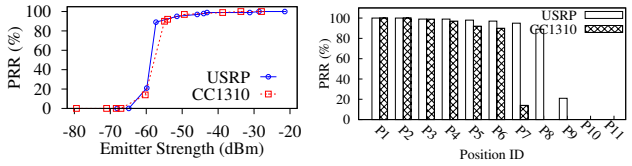
The emitter is placed approximately 15 m from Hark and transmits the injection-locking carrier at 14 dBm using a USRP B200. A COTS SX1262 transmitter sends standard 16-byte LoRa frames on SF-12. At each position, the transmitter sends 100 frames consistent with previous experiments. Both Hark and the COTS SX1262 receiver are placed at the same location at the same height and use the same antenna model [13]. The transmitter is moved across positions P1 through P11 as marked in Fig. 8c. At each position, the COTS receiver reports the RSSI and SNR. We evaluate Hark's performance under all four sampling rates (n, $n/2$, $n/4$, $n/8$). Table 2 reports the results. At +22 dBm, Hark covers most positions in the lab under $n/4$. P9, P10 and P11, separated from the transmitter by multiple concrete partitions, receive no frames at this power.

To push beyond the range limits of the COTS transmitter, we add a power amplifier to reach +30 dBm (1 W), the maximum permitted under FCC regulations for unlicensed spectrum [14]. At this power, previously unreachable positions become viable. P9 jumps from 0% to 100% PRR at Nyquist rate and 95% at $n/4$. P11, which also received no frames at +22 dBm, achieves 96% at Nyquist and 87% at $n/4$. At P10, reception improves from total link failure to 67% at Nyquist. Hark can therefore cover a multi-room laboratory under $n/4$ where detection remains robust even under NLoS.

Outdoor. We also evaluate Hark outdoors under approximate line-of-sight conditions. We pair the COTS SX1262 transmitter with a power amplifier to transmit standard 16-byte LoRa frames on SF-12 at +30 dBm. The emitter is placed similarly to the indoor experiment setup. Running at $n/4$, Hark achieves 93% PRR at a distance of 497.5 m with an average SNR of −8.1 dB.

6.3 Emitter Implications

Impact of Emitter Placement. The experiments thus far kept the emitter at a fixed location. We now vary it and quantify how Hark performs as the incident emitter strength is varied. We fix the LoRa transmitter 8 m from Hark and move only the emitter across P1–P11 in Fig. 8c. Because the transmitter-to-Hark link is unchanged across all positions, PRR variation is attributable solely to the incident emitter strength. At each position, we measure the incident emitter strength at Hark using the BB60C spectrum analyzer, then transmit 100 frames from the COTS LoRa radio and compute the PRR. We repeat this with two emitters: a USRP B200 amplified to +30 dBm through a monopole antenna [13], and a CC1310 LaunchPad [57] transmitting a continuous wave at +26 dBm through its built-in patch antenna. This tests whether an inexpensive COTS radio can serve as the emitter, similar to prior works [23, 61]. Fig. 9a plots our findings against measured emitter strength, and Fig. 9b against each indoor position in the lab. We note two things: first, Hark holds roughly 90% PRR under $n/4$ sampling as the incident emitter strength falls from −21.5 dBm to −57.3 dBm, then drops sharply once it reaches −59.9 dBm. This is consistent with the injection-locking behavior described in Section 2.4



(a) Emitter Strength (b) Position PRR

Figure 9: Hark's PRR under $n/4$. (a) As the emitter moves through P1–P11, the TDO sustains near-90% PRR until the carrier falls to roughly -60 dBm. (b) PRR across indoor positions for both emitters.

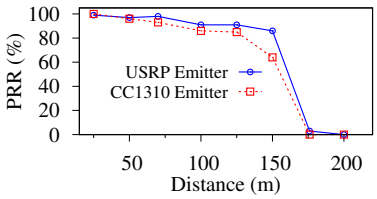


Figure 10: Hark's PRR at $n/4$ as the emitter moves away outdoors under LoS to 200 m. USRP holds above 90% to 125 m and 80% to 150 m; CC1310 holds above 90% to 70 m and 80% to 125 m.

where the TDO remains disciplined while the carrier holds it in lock, and loses lock once the emitter becomes too weak. The threshold is higher than the -70 dBm reported for 2-FSK in SoMix [37], since LoRa demodulation requires greater carrier stability. Second, the CC1310 emitter shows similar behavior. Hark sustains near-90% PRR before dropping sharply at -60.3 dBm. Both emitters lose lock within 0.5 dB of each other, which indicates the threshold is set by the TDO rather than the emitter. A commodity radio can therefore discipline Hark. Consequently, a single inexpensive emitter could serve many receivers across a deployment.

Fig. 9b reports the corresponding PRR at each emitter position. The receiver–emitter separation ranges from 9 m at P1 to 44 m at P11 across this multi-walled NLoS environment. At P6 (27 m), both emitters maintain at least 90% PRR. The +30 dBm USRP emitter further achieves 95% PRR at P7 (29 m) and 89% at P8 (28 m), whereas the +26 dBm CC1310 drops sharply to 14% at P7 and does not sustain reception at the remaining positions. Performance is not monotonic with distance: at P9 (25 m), the USRP achieves only 21% PRR, below P7 and P8 despite the shorter separation, because more concrete walls obstruct it. Taking the farthest position at 90% PRR as a conservative estimate, the indoor NLoS coverage is approximately 29 m for the USRP emitter and 27 m for the CC1310 emitter.

We repeat this experiment outdoors under LoS conditions. We again keep the LoRa transmitter and Hark fixed at 8 m and vary the emitter distance till 200 m. Fig. 10 reports the resulting PRR at $n/4$. Both emitters hold above 90% PRR to 70 m. The USRP stays above 90% to 125 m and above 80% to 150 m. The CC1310 stays above 90% to 70 m and above 80% to 125 m. A single emitter can therefore discipline Hark up to 150 m under LoS. Beyond this point, PRR collapses sharply as the TDO loses injection lock.

Regulatory Compliance. Hark falls under the broad spectrum of ultra-low-power systems that require an external carrier [15, 21, 23, 32, 37, 38, 55, 60, 62, 63, 70, 71]. Ultimately, the regulatory compliance of Hark simplifies to the allowance of the emitter signal. In that sense, it is permitted in many jurisdictions including countries like Singapore [22], New Zealand [44], and Australia [4], where the injection-locking carrier falls under non-specific SRD provisions, although the allowed power varies and is limited in some regions.

6.4 SOTA Comparison

Thus far, we have stated the energy consumption of SOTA receivers in isolation. But receivers like Hark and Sisyphus target large-scale deployments and rely on external assisting signals. Therefore, we quantify their energy and spectral efficiency from an overall network perspective. To this end, we simulate a 500-node network communicating via SF-12 on a standard eight-channel LoRa deployment. We evaluate three representative systems: Hark, Sisyphus, and the COTS SX1262. We vary the network load in steps to compute the average network energy per delivered frame. For simplicity, we assume all transmitted frames are successfully received. In a COTS-based network, the per-frame energy is simply the sum of the node's transmitter and receiver energy. Sisyphus requires a second LoRa frame on an adjacent channel as a helper signal along with each transmitted frame, whereas Hark amortizes the energy of a single 1 W emitter across all frames. Consequently, Sisyphus can only use half of the available spectrum for data communication. Both the COTS network and Hark use the full spectrum, since Hark's narrow emitter can sit outside the data bands without disrupting communication. We assume a standard 1318.9 ms time-on-air for a 16-byte SF-12 frame. Since all three receivers listen continuously, each node's hourly receive energy is divided across the frames within that hour. For the transmitter, we employ the COTS SX1262 at +22 dBm and use its reported power consumption of 389.4 mW required to transmit a 22 dBm frame [51].

We increase the load until all eight channels saturate, and plot the results in Fig. 11. Owing to its ultra-low 162.7 µW operation, Sisyphus consumes substantially less energy per frame than Hark throughout its operating range. However, its spectral utilization terminates at half the maximum load of Hark and COTS. Furthermore, the energy gap narrows as load increases, dropping to under 10% by the time Sisyphus reaches its ceiling. While Hark does trade energy efficiency for higher range and COTS compatibility, it reaches roughly 500 m compared to Sisyphus's 73 m, alongside a complete end-to-end demodulation and decoding pipeline.

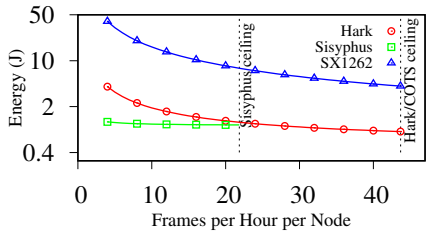


Figure 11: Energy per frame across a simulated network of 500 nodes for increasing loads until all eight channels are saturated.

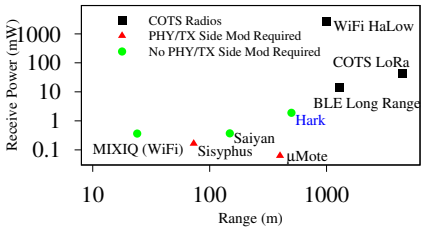


Figure 12: Receive-power versus range tradeoff across low-power receiver architectures. Hark, Sisyphus, and MIXIQ additionally require a receiver-side carrier, which Hark amortizes across many receivers.

Fig. 12 maps Hark against commercial baselines and other low-power receivers by range and power.

Trade-offs. Hark makes select compromises. Hark uses LoRa modulation, which inherently trades off data rate to gain longer range. As such, Hark cannot reach the high data rates that BLE Long Range or Wi-Fi HaLow can offer. Hark's second tradeoff concerns LoRa's SNR. The COTS counterpart can demodulate SF-12 −20 dB below the noise floor. Under $n/4$, Hark holds 100% PRR to −6.5 dB and continues to deliver frames down to −14.5 dB, covering roughly three quarters of this budget but at reduced reliability. These trade-offs are deliberate, and buy a power reduction of over 22×.

When Hark is the Right Choice. Each of these designs suits a different deployment. Sisyphus and μMote achieve substantially lower receiver power and are the right choice when transmitter-side modification or coordinated helper transmissions are acceptable and receiver energy is the dominant constraint. Saiyan preserves transmitter compatibility at lower power for shorter-range deployments. Hark targets a complementary point in this design space: deployments where already-installed LoRa transmitters and gateways must remain unmodified and downlink range matters. Because Hark requires no transmitter-side changes, it can receive standard LoRa downlinks from existing gateway chipsets [48, 50]. This allows integrating with an existing ecosystem of over 125 million deployed devices [35] and their standard LoRaWAN application servers to push on-demand downlinks. Hark retains LoRa's dechirping gain, achieves −62 dBm sensitivity, and reaches 497.5 m. It pays for this with 1.90 mW and a shared receiver-side emitter effective to ∼150 m under LoS.

Table 3: Power breakdown of Hark by operating mode.

Component	Detection	Demodulation
TDO Frontend	0.06 mW	0.06 mW
BPF + Amplifier (LMV641)	1.38 mW	1.38 mW
ADC (ADS7042)	0.04 mW	0.13 mW
MCU (Apollo4 Plus)	0.42 mW	1.41 mW
Total	**1.90 mW**	**2.98 mW**

6.5 Power Breakdown

We measure the power consumption of each component in the Hark prototype across both operating modes. During detection, the MCU runs at 48 MHz and drives the ADS7042 ADC at 62.5 kSPS (under $n/4$ operation). Upon confirming a frame, Hark switches to demodulation mode. The MCU clock is increased to 96 MHz and the ADC sampling rate is also increased to 500 kSPS (2× Nyquist) for improved demodulation performance. The TDO frontend, bandpass filter, and three-stage LMV641 amplifier run without change during both detection and demodulation modes.

Table 3 reports the measured power draw of each component. The analog chain dominates the power budget at 1.44 mW, of which the amplifier and filter account for 1.38 mW. The TDO frontend contributes only 60 µW. During detection, the digital subsystem adds 0.46 mW, bringing total idle listening power to 1.90 mW. During demodulation, the digital subsystem rises to 1.54 mW due to the higher MCU clock and ADC sampling rate which brings the total to 2.98 mW.

7 CONCLUSIONS

This work presented Hark, an ultra-low-power LoRa receiver that enables continuous listening at 1.90 mW while maintaining COTS compatibility. Hark is the first receiver to implement the complete LoRa demodulation and decoding pipeline at this power level. It receives standard LoRa frames from any unmodified COTS transmitter and decodes them through the full chain: synchronization, dechirping, Gray decoding, deinterleaving, Hamming decoding, and CRC verification. The entire pipeline runs in C on an ultra-low-power MCU using the ARM CMSIS library. Hark reduces LoRa receive power by over 22× compared to the SX1262. Our outdoor experiments confirm 93% PRR at nearly 500 m at −8.1 dB SNR. Indoor experiments cover a multi-room NLoS environment through concrete walls. These results bring LoRaWAN Class C within practical reach for battery-powered devices and open the door to on-demand downlinks in deployments that today are limited to Class A.

Acknowledgments

We sincerely thank the reviewers and shepherd for valuable feedback. This work was supported by MOE Tier-1 grant (A-8001661-00-00), ODPRT startup grant (A-8000277-00-00), and Google Research Scholar (A-8002307-00-00-00) and Google TPU Builders (A-8004887-00-00) grants.

References

[1] Francesco Amato. 2017. *Achieving hundreds-meter ranges in low powered RFID systems with quantum tunneling tags.* Ph. D. Dissertation. GaTecH.

[2] Francesco Amato, Christopher W. Peterson, Muhammad B. Akbar, and Gregory D. Durgin. 2015. Long range and low powered RFID tags with tunnel diode. In *2015 IEEE International Conference on RFID Technology and Applications (RFID-TA)*. 182–187. doi:10.1109/RFID-TA.2015.7379815

[3] Francesco Amato, Christopher W. Peterson, Brian P. Degnan, and Gregory D. Durgin. 2018. Tunneling RFID Tags for Long-Range and Low-Power Microwave Applications. *IEEE Journal of Radio Frequency Identification* 2, 2 (2018), 93–103. doi:10.1109/JRFID.2018.2852498

[4] Australian Communications and Media Authority. 2025. Radiocommunications (Low Interference Potential Devices) Class Licence 2025. Federal Register of Legislation, F2025L01047. Registered 5 September 2025; commenced 1 October 2025. https://www.legislation.gov.au/F2025L01047. Accessed: 2026-08-10.

[5] Ravneet Bajwa, Ram Rajagopal, Erdem Coleri, Pravin Varaiya, and Christopher Flores. 2013. In-pavement wireless Weigh-In-Motion. In *2013 ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*. 103–114. doi:10.1145/2461381.2461397

[6] Elizabeth A. Basha, Sai Ravela, and Daniela Rus. 2008. Model-based monitoring for early warning flood detection. In *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems* (Raleigh, NC, USA) *(SenSys '08)*. Association for Computing Machinery, New York, NY, USA, 295–308. doi:10.1145/1460412.1460442

[7] Scott G. Burman, Jingya Gao, Gregory B. Pasternack, Nann A. Fangue, Paul Cadrett, Elizabeth Campbell, and Dipak Ghosal. 2022. TempMesh – A Flexible Wireless Sensor Network for Monitoring River Temperatures. *ACM Trans. Sen. Netw.* 19, 1, Article 15 (Dec. 2022), 28 pages. doi:10.1145/3542697

[8] Yuen Hui Chee, Jan M. Rabaey, and Ali Niknejad. 2006. *Ultra Low Power Transmitters for Wireless Sensor Networks.* Technical Report UCB/EECS-2006-57. Electrical Engineering and Computer Sciences, University of California at Berkeley. http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-57.html

[9] Tuochao Chen, Justin Chan, and Shyamnath Gollakota. 2022. Underwater messaging using mobile devices. In *Proceedings of the ACM SIGCOMM 2022 Conference* (Amsterdam, Netherlands) *(SIGCOMM '22)*. Association for Computing Machinery, New York, NY, USA, 545–559. doi:10.1145/3544216.3544258

[10] Xingda Chen, Deepak Ganesan, Jeremy Gummeson, and Mohammad Rostami. 2021. Cocoon: A Conductive Substrate-based Coupled Oscillator Network for Wireless Communication. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems (SenSys '21)*. Association for Computing Machinery, New York, NY, USA, 84–96.

[11] Rajashekar Reddy Chinthalapani, Dhairya Jigar Shah, Nobel Ang, and Ambuj Varshney. 2025. AudioCast: Enabling Ubiquitous Connectivity for Embedded Systems through Audio-broadcasting Low-power Tags. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 9, 2, Article 27 (June 2025), 32 pages. doi:10.1145/3729471

[12] Leo Esaki, Yasuhiko Arakawa, and Masatoshi Kitamura. 2010. Esaki diode is still a radio star, half a century on. *Nature* 464, 7285 (March 2010), 31. doi:10.1038/464031b

[13] Ettus Research. [n. d.]. VERT900 Vertical Antenna. https://www.ettus.com/all-products/vert900/. Accessed: 2024-12-01.

[14] Federal Communications Commission. 2025. Title 47 CFR Part 15 - Radio Frequency Devices. https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-15 Accessed: March 26, 2025.

[15] Ander Galisteo, Ambuj Varshney, and Domenico Giustiniano. 2020. Two to Tango: Hybrid Light and Backscatter Networks for next Billion Devices. In *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services* (Toronto, Ontario, Canada) *(MobiSys '20)*. Association for Computing Machinery, New York, NY, USA, 80–93. doi:10.1145/3386901.3388918

[16] Amalinda Gamage, Jansen Liando, Chaojie Gu, Rui Tan, Mo Li, and Olivier Seller. 2023. LMAC: Efficient Carrier-Sense Multiple Access for LoRa. *ACM Trans. Sen. Netw.* 19, 2 (2023). doi:10.1145/3564530

[17] Amalinda Gamage, Pramuka Sooriya Patabandige, Spanddhana Sara, and Ambuj Varshney. 2026. Hark Source Code. https://nws.sg/amalinda/hark. Software Repository.

[18] Xiuzhen Guo, Longfei Shangguan, Yuan He, Nan Jing, Jiacheng Zhang, Haotian Jiang, and Yunhao Liu. 2022. Saiyan: Design and Implementation of a Low-power Demodulator for LoRa Backscatter Systems. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 437–451. https://www.usenix.org/conference/nsdi22/presentation/guo

[19] Xiuzhen Guo, Longfei Shangguan, Yuan He, Jia Zhang, Haotian Jiang, Awais Ahmad Siddiqi, and Yunhao Liu. 2020. Aloba: rethinking ON-OFF keying modulation for ambient LoRa backscatter. In *Proceedings of the 18th Conference on Embedded Networked Sensor Systems* (Virtual Event, Japan) *(SenSys '20)*. Association for Computing Machinery, New York, NY, USA, 192–204. doi:10.1145/3384419.3430719

[20] Mehrdad Hessar, Ali Najafi, and Shyamnath Gollakota. 2019. NetScatter: Enabling Large-Scale Backscatter Networks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 271–284. https://www.usenix.org/conference/nsdi19/presentation/hessar

[21] Pan Hu, Pengyu Zhang, Mohammad Rostami, and Deepak Ganesan. 2016. Braidio: An Integrated Active-Passive Radio for Mobile Devices with Asymmetric Energy Budgets. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) *(SIGCOMM '16)*. Association for Computing Machinery, New York, NY, USA, 384–397. doi:10.1145/2934872.2934902

[22] Infocomm Media Development Authority of Singapore. 2023. *Technical Specification for Short Range Devices (SRD).* Technical Report IMDA TS SRD, Issue 1 Revision 3. Infocomm Media Development Authority (IMDA), Telecommunications Standards Advisory Committee (TSAC). http://www.imda.gov.sg

[23] Vikram Iyer, Vamsi Talla, Bryce Kellogg, Shyamnath Gollakota, and Joshua Smith. 2016. Inter-Technology Backscatter: Towards Internet Connectivity for Implanted Devices. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) *(SIGCOMM '16)*. Association for Computing Machinery, New York, NY, USA, 356–369. doi:10.1145/2934872.2934894

[24] Aakriti Jain, Venkata Prashant Modekurthy, and Abusayeed Saifullah. 2026. Real-Time Scheduling and Control Co-Design over LPWAN. *ACM Trans. Cyber-Phys. Syst.* 10, 3, Article 27 (May 2026), 26 pages. doi:10.1145/3785136

[25] Giannis Kazdaridis, Nikos Sidiropoulos, Ioannis Zografopoulos, and Thanasis Korakis. 2021. eWake: A Novel Architecture for Semi-Active Wake-Up Radios Attaining Ultra-High Sensitivity at Extremely-Low Consumption. arXiv:2103.15969 [cs.AR] https://arxiv.org/abs/2103.15969

[26] HJ. Kim, K. A. Sudduth, and J. W. Hummel. 2009. Soil macronutrient sensing for precision agriculture. *Journal of Environmental Monitoring* 11, 10 (2009), 1810–1824.

[27] W. Ko, W. Thompson, and E. Yon. 1962. A tunnel-diode FM transmitter for medical research and laboratory telemetering. In *1962 IEEE International Solid-State Circuits Conference. Digest of Technical Papers*, Vol. V. 70–71. doi:10.1109/ISSCC.1962.1157367

[28] Lakshman Krishnamurthy, Robert Adler, Phil Buonadonna, Jasmeet Chhabra, Mick Flanigan, Nandakishore Kushalnagar, Lama Nachman,

and Mark Yarvis. 2005. Design and deployment of industrial sensor networks: experiences from a semiconductor plant and the north sea. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems* (San Diego, California, USA) *(SenSys '05)*. Association for Computing Machinery, New York, NY, USA, 64–75. doi:10.1145/1098918.1098926

[29] Andreas Kuster, Huatao Xu, Rui Tan, and Mo Li. 2026. RF Super Resolution: A Deep Learning Approach to Spatial Enhancement for LoRa. In *Proceedings of the 24th Annual International Conference on Mobile Systems, Applications and Services* (University of Cambridge, Cambridge, United Kingdom) *(MobiSys '26)*. Association for Computing Machinery, New York, NY, USA, 464–477. doi:10.1145/3745756.3809216

[30] In Lee. 2016. An exploratory study of the impact of the internet of things IoT on business model innovation: Building smart enterprises at fortune 500 companies. *Int. J. Inf. Syst. Soc. Chang.* 7, 3 (Jul 2016), 1–15.

[31] Jansen C. Liando, Amalinda Gamage, Agustinus W. Tengourtius, and Mo Li. 2019. Known and Unknown Facts of LoRa: Experiences from a Large-scale Measurement Study. *ACM Trans. Sen. Netw.* 15, 2, Article 16 (Feb. 2019), 35 pages. doi:10.1145/3293534

[32] Vincent Liu, Aaron Parks, Vamsi Talla, Shyamnath Gollakota, David Wetherall, and Joshua R. Smith. 2013. Ambient backscatter: wireless communication out of thin air. In *Proceedings of the ACM SIGCOMM 2013 Conference* (Hong Kong, China) *(SIGCOMM '13)*. Association for Computing Machinery, New York, NY, USA, 39–50. doi:10.1145/2486001.2486015

[33] Wenquan Liu, Xin Wang, Yongxin Song, Ruirui Cao, Liangliang Wang, Zhengguang Yan, and Guiye Shan. 2020. Self-powered forest fire alarm system based on impedance matching effect between triboelectric nanogenerator and thermosensitive sensor. *Nano Energy* 73 (2020), 104843. doi:10.1016/j.nanoen.2020.104843

[34] LoRa Alliance. 2018. *LoRaWAN® Specification v1.0.3*. Technical Specification. LoRa Alliance. Available at https://resources.lora-alliance.org/technical-specifications/lorawan-specification-v1-0-3.

[35] LoRa Alliance. 2025. *LoRa Alliance Reports 125 Million LoRaWAN End Devices Deployed Globally*. https://resources.lora-alliance.org/home/lora-alliance-reports-125-million-lorawan-end-devices-deployed-globally

[36] Jiakang Lu, Tamim Sookoor, Vijay Srinivasan, Ge Gao, Brian Holben, John Stankovic, Eric Field, and Kamin Whitehouse. 2010. The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems* (Zürich, Switzerland) *(SenSys '10)*. Association for Computing Machinery, New York, NY, USA, 211–224. doi:10.1145/1869983.1870005

[37] Sooriya Patabandige Pramuka Medaranga, Rajashekar Reddy Chinthalapani, Wenqing Yan, Prabal Dutta, and Ambuj Varshney. 2025. Unraveling the Missing Link in Low-power Communication: An Autodyning Receiver Architecture that Achieves a Long Range. In *The 23rd Annual International Conference on Mobile Systems, Applications and Services (MobiSys '25)*. Association for Computing Machinery, Anaheim, CA, USA, 1–14. doi:10.1145/3711875.3729164

[38] Muhammad Sarmad Mir, Wenqing Yan, Prabal Dutta, Domenico Giustiniano, and Ambuj Varshney. 2023. TunnelLiFi: Bringing LiFi to Commodity Internet of Things Devices. In *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications* (Newport Beach, California) *(HotMobile '23)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3572864.3580327

[39] Yamir Moreno and Amalio F. Pacheco. 2004. Synchronization of Kuramoto Oscillators in Scale-Free Networks. *EPL (Europhysics Letters)* 68, 4 (2004), 603.

[40] Nordic Semiconductor. 2025. *Power Profiler Kit II: Power profiling and power optimization of embedded solutions*. https://www.nordicsemi.com/Products/Development-hardware/Power-Profiler-Kit-2

[41] P. O'Donovan, K. Leahy, K. Bruton, and D. T. J. O'Sullivan. 2015. An industrial big data pipeline for data-driven analytics maintenance applications in large-scale smart manufacturing facilities. *Journal of Big Data* 2, 1 (Nov 2015), 25. doi:10.1186/s40537-015-0034-z

[42] Yao Peng, Vamsi Talla, and Shyam Gollakota. 2018. PLoRa: A Passive Long-range Data Network from Ambient LoRa Transmissions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*. ACM, 147–160.

[43] Rajeev Piyare, Amy L. Murphy, Michele Magno, and Luca Benini. 2018. On-Demand TDMA for Energy Efficient Data Collection with LoRa and Wake-up Receiver. In *2018 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. 1–4. doi:10.1109/WiMOB.2018.8589151

[44] Radio Spectrum Management, Ministry of Business, Innovation and Employment. 2022. Radiocommunications Regulations (General User Radio Licence for Short Range Devices) Notice 2022. New Zealand Gazette, Notice No. 2022-go3100; Licence No. 296531. In force 1 August 2022. https://gazette.govt.nz/notice/id/2022-go3100. Accessed: 2026-08-10.

[45] Pieter Robyns, Peter Quax, Wim Lamotte, and William Thenaers. 2017. gr-lora: An efficient LoRa decoder for GNU Radio. *Zenodo ed* 10 (2017), 5281.

[46] Mohammad Rostami, Xingda Chen, Yuda Feng, Karthikeyan Sundaresan, and Deepak Ganesan. 2021. MIXIQ: re-thinking ultra-low power receiver design for next-generation on-body applications. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking* (New Orleans, Louisiana) *(MobiCom '21)*. Association for Computing Machinery, New York, NY, USA, 364–377. doi:10.1145/3447993.3483270

[47] RCA Corporation. Semiconductor and Materials Division. 1963. *RCA Tunnel Diode Manual*. RCA. https://books.google.se/books?id=2G5_ugEACAAJ

[48] Semtech Corporation. 2017. SX1301 Datasheet. Datasheet, Rev. 2.4. https://www.semtech.com/products/wireless-rf/lora-core/sx1301, accessed 2026-08-04.

[49] Semtech Corporation. 2019. *SX1276/77/78/79 - LoRa Modem Datasheet*. Semtech Corporation. https://www.semtech.com/products/wireless-rf/lora-transceivers/sx1276.

[50] Semtech Corporation. 2020. SX1302 Datasheet. Datasheet, Rev. 1.2. https://www.semtech.com/products/wireless-rf/lora-core/sx1302, accessed 2026-08-04.

[51] Semtech Corporation. 2023. *SX1261/62/68 - LoRa Modem Datasheet*. Semtech Corporation. https://www.semtech.com/products/wireless-rf/lora-transceivers/sx1262.

[52] Semtech Corporation. 2025. *LR1121 Datasheet*. Semtech Corporation. https://www.semtech.com/products/wireless-rf/lora-connect/lr1121.

[53] Emiliano Sisinni, Abusayeed Saifullah, Song Han, Ulf Jennehag, and Mikael Gidlund. 2018. Industrial internet of things: Challenges, opportunities, and directions. *IEEE Transactions on Industrial Informatics* 14, 11 (2018), 4724–4734.

[54] Yihang Song, Li Lu, Jiliang Wang, Chong Zhang, Hui Zheng, Shen Yang, Jinsong Han, and Jian Li. 2023. μMote: Enabling Passive Chirp De-spreading and uW-level Long-Range Downlink for Backscatter Devices. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 1751–1766. https://www.usenix.org/conference/nsdi23/presentation/song-yihang

[55] Vamsi Talla, Mehrdad Hessar, Bryce Kellogg, Ali Najafi, Joshua R. Smith, and Shyamnath Gollakota. 2017. LoRa Backscatter: Enabling The Vision of Ubiquitous Connectivity. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 3, Article 105 (Sept. 2017), 24 pages.

doi:10.1145/3130970

[56] Joachim Tapparel, Orion Afisiadis, Paul Mayoraz, Alexios Balatsoukas-Stimming, and Andreas Burg. 2020. An Open-Source LoRa Physical Layer Prototype on GNU Radio. In *2020 IEEE 21st International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. 1–5. doi:10.1109/SPAWC48557.2020.9154273

[57] Texas Instruments. [n. d.]. CC1310 Sub-1 GHz Wireless Microcontroller. https://www.ti.com/product/CC1310. Accessed: 2024-12-09.

[58] Texas Instruments. 2025. ADS7042 Datasheet. https://www.ti.com/lit/ds/symlink/ads7042.pdf?ts=1761807058222. Accessed: 2025-11-28.

[59] Harry L. Van Trees. 2001. *Detection, Estimation, and Modulation Theory, Part I: Detection, Estimation, and Linear Modulation Theory.* John Wiley & Sons, New York.

[60] Ambuj Varshney and Lorenzo Corneo. 2020. Tunnel emitter: tunnel diode based low-power carrier emitters for backscatter tags. In *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking* (London, United Kingdom) *(MobiCom '20)*. Association for Computing Machinery, New York, NY, USA, Article 42, 14 pages. doi:10.1145/3372224.3419199

[61] Ambuj Varshney, Oliver Harms, Carlos Pérez-Penichet, Christian Rohner, Frederik Hermans, and Thiemo Voigt. 2017. LoRea: A Backscatter Architecture that Achieves a Long Communication Range. In *Proceedings of the 15th ACM Conference on Embedded Network Sensor Systems* (Delft, Netherlands) *(SenSys '17)*. Association for Computing Machinery, New York, NY, USA, Article 18, 14 pages. doi:10.1145/3131672.3131691

[62] Ambuj Varshney, Andreas Soleiman, and Thiemo Voigt. 2019. TunnelScatter: Low Power Communication for Sensor Tags using Tunnel Diodes. In *The 25th Annual International Conference on Mobile Computing and Networking* (Los Cabos, Mexico) *(MobiCom '19)*. Association for Computing Machinery, New York, NY, USA, Article 50, 17 pages. doi:10.1145/3300061.3345451

[63] Ambuj Varshney, Wenqing Yan, and Prabal Dutta. 2022. Judo: addressing the energy asymmetry of wireless embedded systems through tunnel diode based wireless transmitters. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services* (Portland, Oregon) *(MobiSys '22)*. Association for Computing Machinery, New York, NY, USA, 273–286. doi:10.1145/3498361.3538923

[64] Deepak Vasisht, Zerina Kapetanovic, Jongho Won, Xinxin Jin, Ranveer Chandra, Sudipta Sinha, Ashish Kapoor, Madhusudhan Sudarshan, and Sean Stratman. 2017. FarmBeats: An IoT Platform for Data-Driven Agriculture. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA, 515–529. https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/vasisht

[65] Han Wang, Yihang Song, Qianhe Meng, Zetao Gao, Chong Zhang, and Li Lu. 2024. Sisyphus: Redefining Low Power for LoRa Receiver. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking* (Washington D.C., DC, USA) *(MobiCom '24)*. Association for Computing Machinery, New York, NY, USA, 1177–1191. doi:10.1145/3636534.3690686

[66] Xieyang Wu, Rajalakshmi Nandakumar, and Shyam Gollakota. 2021. P2LoRa: A Passive Long-range Data Network from Ambient LoRa Transmissions. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking (MobiCom '21)*. Association for Computing Machinery, 1–14.

[67] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. FTrack: Parallel Decoding for LoRa Transmissions. *IEEE/ACM Transactions on Networking* 28, 6 (2020), 2573–2586. doi:10.1109/TNET.2020.3018020

[68] Xianjin Xia, Yuanqing Zheng, and Tao Gu. 2020. LiteNap: Downclocking LoRa Reception. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 2321–2330. doi:10.1109/INFOCOM41043.2020.9155224

[69] Zhenqiang Xu, Shuai Tong, Pengjin Xie, and Jiliang Wang. 2023. From Demodulation to Decoding: Toward Complete LoRa PHY Understanding and Implementation. *ACM Trans. Sen. Netw.* 18, 4, Article 64 (Jan. 2023), 27 pages. doi:10.1145/3546869

[70] Pengyu Zhang, Pan Hu, Vijay Pasikanti, and Deepak Ganesan. 2014. EkhoNet: High Speed Ultra Low-Power Backscatter for next Generation Sensors. In *Proceedings of the 20th Annual International Conference on Mobile Computing and Networking* (Maui, Hawaii, USA) *(MobiCom '14)*. Association for Computing Machinery, New York, NY, USA, 557–568. doi:10.1145/2639108.2639138

[71] Pengyu Zhang, Mohammad Rostami, Pan Hu, and Deepak Ganesan. 2016. Enabling Practical Backscatter Communication for On-body Sensors. In *Proceedings of the 2016 ACM SIGCOMM Conference* (Florianopolis, Brazil) *(SIGCOMM '16)*. Association for Computing Machinery, New York, NY, USA, 370–383. doi:10.1145/2934872.2934901

[72] Xinyu Zhang and Kang G. Shin. 2012. E-MiLi: Energy-Minimizing Idle Listening in Wireless Networks. *IEEE Transactions on Mobile Computing* 11, 9 (2012), 1441–1454. doi:10.1109/TMC.2012.112