Echoes of Earth: Building an Autonomous Environmental Lab for Acoustic Sensing

Hudson Reynolds*
Boston University
Boston, USA
hudsonre@bu.edu

Alex Tuecke*
Worcester Polytechnic Institute
Worcester, USA
ahtuecke@wpi.edu

Mike Sherman (advisor)
University of Chicago
Chicago, Illinois, USA
shermanm@uchicago.edu

Kate Keahey (advisor)
Argonne National Laboratory
Lemont, Illinois, USA
keahey@uchicago.edu

ABSTRACT

There is a critical need to provide natural resource managers with real-time bioacoustic information for biodiversity conservation. We develop a prototype system for real-time, inexpensive, high-quality soundscaping at scale. Our system eliminates the need for manual data retrieval and expert maintenance, drastically reducing operational costs. We create Listener, a solar-powered recording and streaming device built with ESP32 and AudioMoth, and Aggregator, based on Raspberry Pi 5, to collect streams from multiple Listeners over WiFi HaLow and perform local inference to analyze recordings. Aggregators upload the collected raw data, analysis results, and performance metrics to the Chameleon testbed. We present a live data visualization Dashboard for metrics, analysis results, and an object store for raw recordings. We show how the system enables novel land management techniques by deploying at organic vineyards in Michigan. We inspect the power consumption, cost, and capabilities of our system. We achieve success in key metrics and show scalability to 25 concurrent Listeners per Aggregator.

CCS CONCEPTS

• **Computer systems organization → Sensor networks**; **Real-time system architecture**; *Cloud computing*; • **Hardware →** *System-level fault tolerance*; • **Applied computing → Agriculture**; **Environmental sciences**.

KEYWORDS

real-time, distributed, sensing, autonomous, solar, halow, birdnet

ACM Reference Format:
Hudson Reynolds, Alex Tuecke, Mike Sherman (advisor), and Kate Keahey (advisor). 2025. Echoes of Earth: Building an Autonomous Environmental Lab for Acoustic Sensing. In *Proceedings of ACM Student Poster Competition (ACM '25)*. ACM, New York, NY, USA, 3 pages. https://doi.org/XXXXXXX.XXXXXXX

*Both authors contributed equally to this research.

1 INTRODUCTION

The growth of bioacoustic monitoring has outpaced the capabilities of existing recording technologies, creating significant barriers to large-scale, real-time monitoring. The industry-standard recording devices, Wildlife Acoustics devices, cost $600-$1000+.[1] They require manual data retrieval, preventing real-time insights, and depend on expert setup and maintenance, which drives up operational costs and limits scalability. This effort limits the amount of data that can be collected and processed, and the areas and phenomena that can be studied. While devices are deployed, there is no status or heartbeat, which can lead to data loss if a malfunctioning device is not identified for months on end. As deployments scale, the lack of a management interface means increasing organization overhead. We present an end-to-end solution for alleviating these burdens.

2 APPROACH

We chose WiFi HaLow (802.11ah) as the Listener to Aggregator link because it supports up to 32.5 Mbps and over 1 KM of range. LoRa, Zigbee, and Meshtastic do not support this throughput, and other alternatives like satellite are either too expensive in terms of power or cost, or require infrastructure like cell towers. [2] [3]

2.1 Listener

Listener uses an ESP32S3, a WiFi HaLow transceiver, a 512GB MicroSD, and a 10 watt Voltaic Solar Panel Battery combo for power. We use the open-source AudioMoth Dev running AudioMoth-USB-Microphone firmware. We package Listener into the weatherproof AudioMoth Dev Enclosure, bringing the total Listener cost to $375. Listener streams 16-bit PCM data at 48 KHz (768 Kbps) to match the quality of recordings that field ecologists commonly use. Listener streams to an HTTP endpoint hosted by Aggregator on the HaLow network. If the network is down or busy, Listener uses the onboard MicroSD card to buffer up to 8 weeks of continuous data and catches up seamlessly when the network returns.

2.2 Aggregator

Aggregator is a Raspberry Pi 5 plugged into a Heltec 7608 WiFi HaLow Router via ethernet, with a 2.4GHz WiFi uplink to Chameleon.

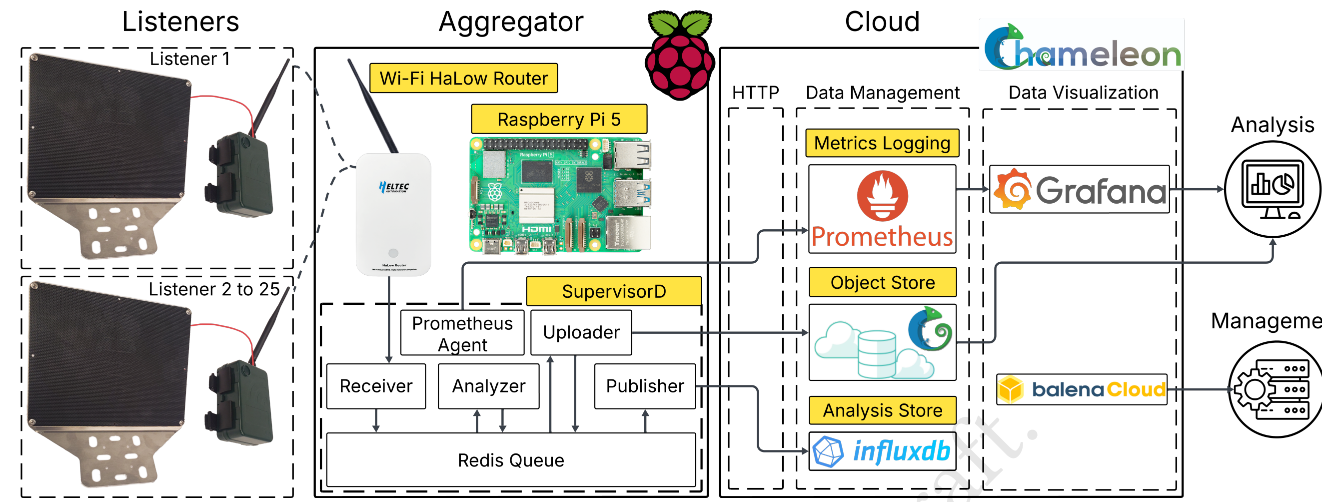

Figure 1: Architecture Diagram

(a) Upload (b) Analyze

(c) Analyze and Upload (d) Redis Job Queue Depth

Figure 2: CPU Utilization over time (a) Upload recordings, (b) Analyze with BirdNET, (c) Simultaneous analyze and upload, and (d) Redis job queue depth over time, for 1, 5, and 25 simulated Listeners streaming to a real Aggregator

This uplink can be whatever the specific deployment location supports. This extended star architecture reduces the infrastructure required to deploy at scale. Aggregator runs BalenaOS for scalable deployments and management through Balena Cloud Fleets. A Dockerized stack manages the Aggregator code modules, which use an in-memory Redis Queue to communicate with each other efficiently, reducing small writes and extending the boot MicroSD longevity. The total cost of Aggregator is $210. Aggregator runs the Cornell BirdNET model locally for analysis.

Table 1: Aggregator Power Draw with 25 Listeners

Task	Power (W)
Upload	3.7
Analyze	5.5
Both	6.1

2.3 Cloud

We host the Grafana Dashboard, Prometheus metrics time series database, and InfluxDB3 analysis database on a Chameleon Kernel-Based Virtual Machine (CHI@KVM) and use the Chameleon Object Store (CHI@TACC) for raw recordings. We utilize Balena Cloud Fleets, CHI@EDGE infrastructure, and our Dashboard to facilitate scalability, maintainability, and organization.

3 EVALUATION

Figure 2 shows the results of Aggregator tests. The most strenuous task, simultaneous uploading and analyzing for 25 Listeners, has an average CPU utilization of 81% with an average job queue depth of 25.8 upload jobs and 25.7 analyze jobs. These results demonstrate Aggregator can sustain 25 concurrent Listeners without falling behind, validating our architecture's scalability. Table 1 shows the Aggregator power draw during the tests from Figure 2. By eliminating the need for manual retrieval and continuous expert oversight, the system reduces both field labor and downtime risk. Together with low device cost and scalable aggregation, this enables operational savings that make larger deployments feasible.

4 CONCLUSION

This work successfully addresses the critical need for real-time bioacoustic monitoring in biodiversity conservation with a scalable solar-powered solution. We demonstrate that autonomous, high-quality, real-time bioacoustic monitoring is now feasible at scale, while eliminating the need for manual retrieval and expert maintenance, drastically reducing operation costs. This opens new possibilities for data-driven biodiversity conservation and adaptive ecosystem management. Future work will add Listener remote reconfigurability, and create an app to aid with large-scale deployments.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under award OAC-2150500.

REFERENCES

[1] [n. d.]. Song Meter SM4 vs. Song Meter Mini 2 vs. Song Meter Micro 2. https://www.wildlifeacoustics.com/products/song-meter-sm4-vs-mini2-vs-micro-2

[2] Victor Baños-Gonzalez, M. Shahwaiz Afaqui, Elena Lopez-Aguilera, and Eduard Garcia-Villegas. 2016. IEEE 802.11ah: A Technology to Face the IoT Challenge. *Sensors* 16, 11 (Nov. 2016), 1960. https://doi.org/10.3390/s16111960 Publisher: Multidisciplinary Digital Publishing Institute.

[3] Ben Meadors. 2025. Is LongFast Holding Your Mesh Back? Better LoRa Presets for Bigger Meshtastic Networks | Meshtastic. https://meshtastic.org/blog/why-your-mesh-should-switch-from-longfast/

Received 25 August 2025