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

Hudson Reynolds[1], **Alex Tuecke**[2], Mike Sherman (advisor)[3], Kate Keahey (advisor)[3*]
[1]Boston University, [2] Worcester Polytechnic Institute, [3]University of Chicago, [*]Argonne National Laboratory

Real-Time Soundscaping

- Supports biodiversity conservation:
  - Tracking Bengal tigers
  - Detecting malaria-infected mosquito swarms
  - Protecting organic crops, and more
- Studies how the environment and humans affect each other
- Need lots of data over a long period of time
- Current solutions require manual data collection
  - Even if we detect something, can't take action!
- And it's very expensive:
  - Hardware: $1000 per device + Deployment Cost
  - Operation: Hours to Days of expert's time

Figure 1: Borneo
Source: Purdue CGS

How can we:

- Reliably stream and analyze audio data in **real-time**, not days to weeks later
- From **thousands of Listeners**, not dozens
- Minimize hardware and operating costs for years-long studies
- Enable easy deployment, visualization, storage, and management

Figure 2: Expert maintenance
Source: Purdue CGS

The Problem of Scaling

- Raw audio data is 48 KHz, 16 bit WAV, 768 Kbps per Listener
- Listeners deployed in locations with no infrastructure: no power, no network
- We need to deploy wireless infrastructure that can achieve this data rate and be solar powered
- **Approach:**
  - Split the system into "Listeners" and "Aggregators"
  - Listeners are solar powered: send data to a local Aggregator
  - Aggregators handle dozens of Listeners, more feasible to provide power and uplink to just Aggregators
  - Persist the data into a scalable cloud backend
  - Provide management & visualization dashboard with remote reconfigurability

Figure 3: Bird pest in vineyard
Source: Daniela Somers

Listener and Aggregator

Listener:

- Optimize for low power: < 0.7 W
- AudioMoth as USB-microphone
- ESP32 bridges USB to WiFi HaLow: streams data to Aggregator via HaLow AP
- Circular buffer on MicroSD card handles network disruptions of days to weeks

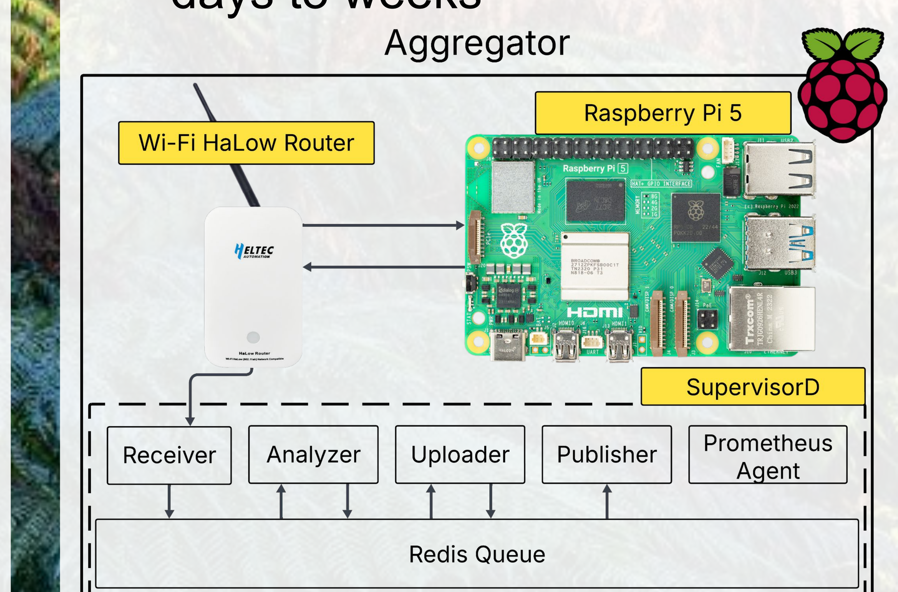

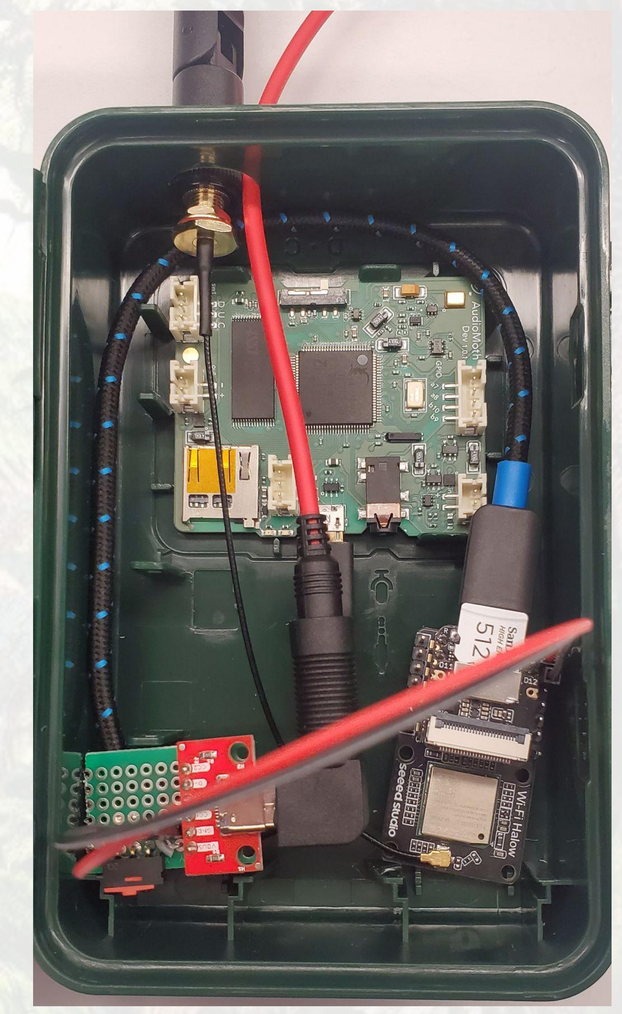

Figures 4 and 5: Internal view of Aggregator (left) and Listener (top)

Services

Each Aggregator runs:

- Data **receiver** and **uploader** services, persists data streams into object storage
- **Analysis** service – runs BirdNET on audio streams, can send alerts locally even if uplink is down
- **Publisher** uploads analysis results to InfluxDB in the cloud, links with audio data in object store
- **Prometheus** exposes hardware and software metrics to monitor the system

Raw data and analysis results are available through the cloud-based Grafana frontend

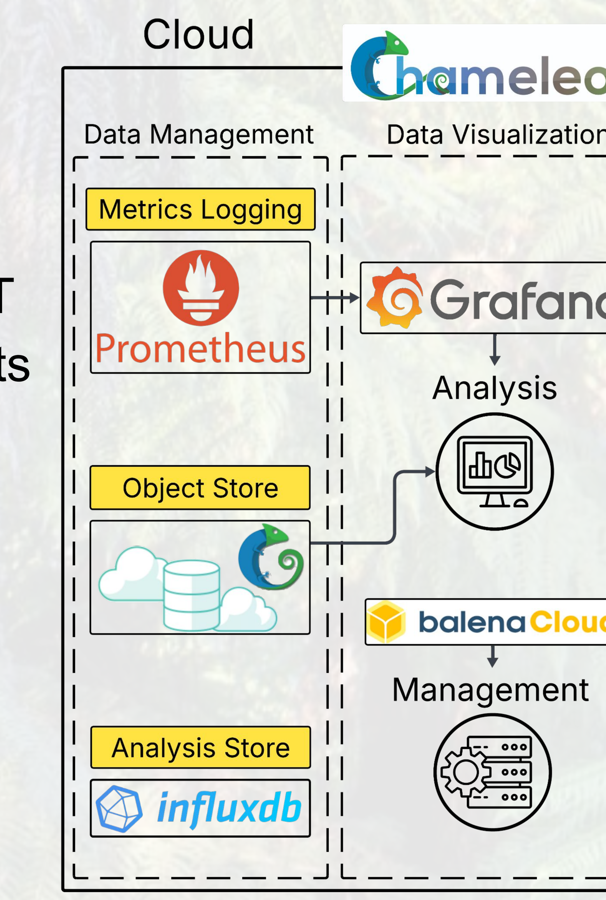

Figure 6: Cloud Services

Results

- ➔ Real-Time: Each Aggregator handles real-time streams of up to 25 Listeners, while running BirdNET
- ➔ Cloud services can be scaled horizontally – increase Listeners by increasing Aggregators
- ➔ Cost minimization:
  - ◆ Benchmark devices costs $1000, prototype Listeners are $375 each, Aggregators are $210 each
  - ◆ Data upload and remote management reduces operational costs
  - ◆ Listener achieves 0.64 W while continuously recording & transmitting
  - ◆ Aggregator averages 6 W while uploading & analyzing 25 Listener streams
- ➔ Lessons learned from ML inference at the edge:
  - ◆ Raspberry Pi 5 achieves 2.5x higher throughput than Pi 4B
  - ◆ BirdNET analysis accuracy: comparable to benchmark device
  - ◆ BirdNET analysis timeliness: Alert comes less than 30 seconds after event

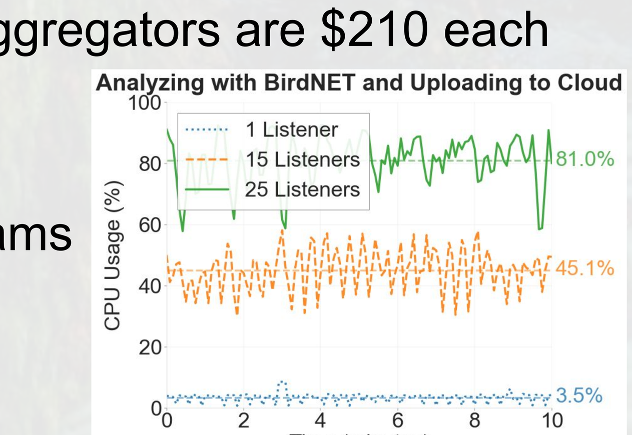

Figure 7: Aggregator CPU for 25 concurrent streams

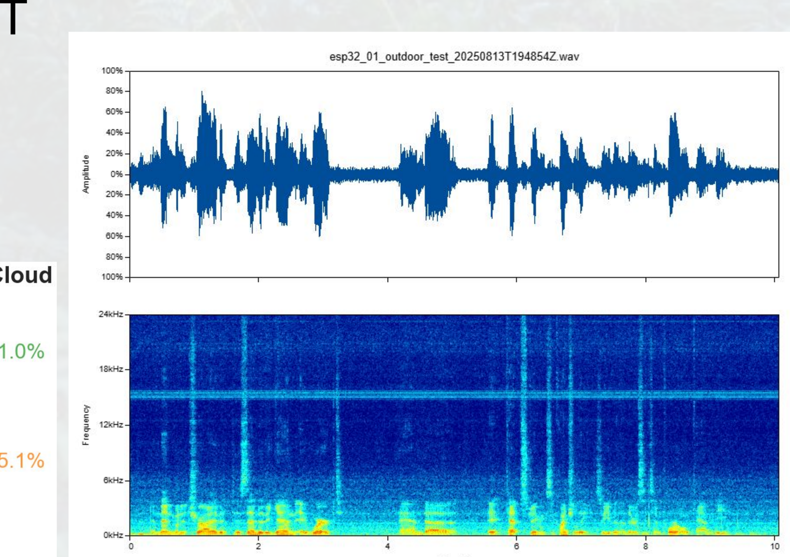

Figure 8: Spectrogram from Listener correctly identified as "Human Speech"

Conclusions

- ➔ Field deployments are vital to understand real vs theoretical performance
- ➔ Listener power usage has lots of room for optimization, but already fits within limit of practical solar panels
- ➔ Already practical for useful deployments: agriculture, pest detection, people detection, etc.
- ➔ Our system eliminates need for most human intervention, and streamlines operational effort which is majority of cost
- ➔ More work needed to make Aggregators practical in remote locations at scale: more solar, satellite, weatherproofing
- ➔ Future work: add Listener remote reconfigurability, reduce Listener power draw, upgrade deployment tools

Figure 9: Link to project website