



Poster: Towards Low-Power Comprehensive Biodiversity Monitoring

Josh Millar, Sarab Sethi, Hamed Haddadi
Imperial College London

Anil Madhavapeddy
University of Cambridge

Abstract

The Kunming-Montreal Global Biodiversity Framework sets ambitious targets for 2030, including halting human-induced species extinction. Achieving these requires comprehensive data on global biodiversity patterns, which can only be gathered through in-situ distributed sensor networks. However, these multi-device networks are constrained by battery lifetimes, must gather rich data from power-hungry sensors, and yet need to be deployed in remote environments for long periods. This note introduces a prototype multi-sensor device, and outlines how embedded scheduling could be used for extending sensor lifetime and resource-efficiency.

CCS Concepts

• **Computing methodologies** → **Reinforcement learning**; • **Computer systems organization** → **Embedded systems**.

Keywords

Edge computing, TinyML, Biodiversity, Low-Power Sensing

ACM Reference Format:

Josh Millar, Sarab Sethi, Hamed Haddadi and Anil Madhavapeddy. 2024. Poster: Towards Low-Power Comprehensive Biodiversity Monitoring. In *The 22nd ACM Conference on Embedded Networked Sensor Systems (SENSYS '24)*, November 4–7, 2024, Hangzhou, China. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3666025.3699400>

1 Introduction

The low-power operation of battery-powered or energy-harvesting sensors is vital for long-term network deployments. This is especially important for biodiversity monitoring networks, which are often deployed in remote, infrastructure-scarce environments, yet must gather rich data from power-hungry sensors (i.e., cameras and microphones). Consequently, biodiversity monitoring devices face a trade-off between data quality and coverage [1, 14, 15]; single-sensor devices [5] are cost-effective and low-power but offer limited data richness, while multi-sensor devices provide richer data but are generally expensive and power-hungry. This note outlines a versatile multi-sensor device supporting low-power operation via adaptive scheduling. The application of biodiversity monitoring presents a unique opportunity here; learning schedules for events of interest (animal vocalizations, feeding activity, etc.) not only optimizes power usage but also generates large-scale information useful for conservationists.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SENSYS '24, November 4–7, 2024, Hangzhou, China

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0697-4/24/11

<https://doi.org/10.1145/3666025.3699400>

Various strategies for adaptive scheduling include context rule-based methods [11], genetic algorithms [3], probabilistic approaches for estimating event likelihoods [7], and reinforcement learning-based schedulers [6]. The latter excel in the absence of predefined event models, especially if deploying many sensors for long durations, as event patterns vary and performance can degrade if scheduling isn't regularly fine-tuned. However, learning-based approaches for network-wide scheduling generally need federated/centralized coordination [10].

We instead look at on-device scheduling for event-driven networks, using low-power collaboration between neighboring devices to minimize redundancy. This note outlines a functional prototype, including camera, microphone, and PIR, and evaluates its operational lifetime using real-world biodiversity data, power measurements, and an embedded, learning-based scheduler. Whilst initially isolated/standalone, its implementation lays the groundwork for a large-scale networked approach.

2 Prototype and Results

Prototype Design. The prototype is built on an ESP32s3-variant¹, and features ultra-low deep-sleep consumption of $\sim 19\mu\text{A}$ (with 3.7V supply). The prototype also includes built-in power management features, such as a LiPo/LiIon battery charger IC and fuel gauge, supporting battery health monitoring and time-to-empty/full battery estimation along with a DC input to source reliable renewable power. The ESP32s3 itself is based on a XTensa SoC, supporting 512KB of internal SRAM and 16KB RTC SRAM for retaining state over deep-sleep. The SoC includes a RISC-V/FSM ultra-low-power coprocessor for interfacing with external I2C sensors while its main processor remains in deep-sleep.

The prototype includes a 5MP camera², an omnidirectional microphone³, and PIR⁴ sensor. These components were selected for their low-power operation and use in related biodiversity monitoring applications [2]. The PIR sensor remains active continuously, even in deep-sleep, acting as an event-trigger for our camera. The microphone is our scheduled component, but recording could also be event-triggered if events are of longer duration.

4G is used for transmitting recordings and images.

We measure the current draw of the device using a high-voltage power monitor capable of μA -scale measurements⁵. The board is supplied at 3.3V via its JST-PH battery connector.

Event Detection and Scheduling. We apply an off-the-shelf bird vocalization detection model, BirdNET [8], to generate events and

¹<https://powerfeather.dev>

²Omnivision OV5640

³INMP441

⁴PaPIR EKMB1103111

⁵<https://www.msoon.com/high-voltage-power-monitor>

durations from continuous audio recorded in the Sabah rainforest, Borneo⁶. We address BirdNET false positives by applying a confidence threshold of 0.7. Users can generate species-specific thresholds based on expert validation of recorded pilot data to further improve detection accuracy [13]. The generated events are used in building and evaluating our scheduler. These, combined with our power measurements, allow us to estimate our prototype’s battery lifetime in a real-world deployment.

We discretize state into t periods of fixed activation rate. This can be done in practice based on observed/expected event patterns and imposed operational constraints. We also use K possible activation frequencies based on the observed/expected range of event intervals and durations. The scheduler uses Q-learning [12], a model-free reinforcement learning approach, given its low memory requirements and $\mathcal{O}(n)$ inference/update [6]. With hourly periods and $K = 7$ possible actions, the resulting 24×7 Q-table is only 3KB, meaning multiple Q-tables can be used on one device for the scheduling of different connected sensors. The reward is structured as a cost-balance between the number of positive (i.e. those associated with detected events) and sensor negative activations.

Q-learning inference and Q-table updates are implemented on-device to show their negligible overhead. We record continuously for 0.1s upon activation to detect events; if an event is detected, recording persists until the event concludes, otherwise the device returns to deep-sleep until its next scheduled activation. Goertzel filtering [16] is used for event detection, processing a 0.1s recording (at 16000kHz) with just .03s latency, minimally impacting our device’s lifetime. However, Goertzel, and other filtering-based approaches detect any event within their range of covered frequencies, while missing off-frequency ones, potentially resulting in degraded fine-tuning performance. We therefore also implement an alternative integer-quantized one-layer convolutional model for event detection, embedded using ESP-TFLite-Micro; its measurement also follows a 0.1s recording. This model is built on spectrograms of 0.1s slices from our Borneo recordings and event labels generated using BirdNET, and reaches $\sim 70\%$ accuracy on unseen event data; however, no validation is done on the positive slices. ESP-DSP is used for extracting spectrograms from buffered recordings on-device.

The schedule is learned over 24 hours of detections, and evaluated over the following 24 hours, to imitate day-by-day learning with minimal pre-deployment data. We compare our scheduler to a number of fixed baselines that activate the device every n seconds, irrespective of the current period. These baselines are commonly used in real-world deployments, alongside continuous periodic recording of fixed duration [5]. More complex algorithmic schedules (e.g. based on known event patterns) could also be implemented; although these are non-adaptable, they can be used for Q-table initialization to accelerate convergence.

Performance and Battery Lifetime. The device’s lifetime is calculated using a 13400mAh LiPo/Li-Ion battery. We model the average operating current based on power measurements of the various device states, including Q-learning/TFLite inference and update, and the number of detected events. We assume that each detected event triggers a camera activation, resulting in a much higher activation rate than observed for bird/mammal detection in the Sabah

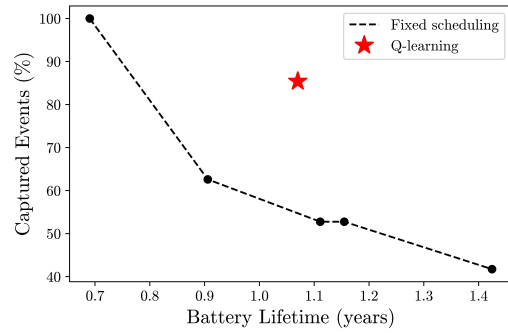


Figure 1: Pareto curve of fixed vs. learned schedules

rainforest [9], followed by a 4G data transmission. We also do not parallelize operations. This gives us an overall relatively conservative estimate of battery lifetime.

The figure above summarizes the scheduling outcomes. We capture a high percentage of events (85.3% vs $\sim 53.5\%$) at much reduced power consumption, extending the device’s lifetime from approximately 0.77 years on a fixed schedule to over a year (1.07 years). Using TFLite for event detection yields similar results (1.05 years). With a renewable power source for recharging, our lifetime bottleneck moves from power consumption to device/network resilience.

3 Next Steps

This note outlines a general event-driven scheduler for isolated devices. However, comprehensive biodiversity monitoring necessitates large-scale deployments of networked sensors, prompting further optimization through coordinated scheduling.

We are now exploring a uniquely decentralized, scalable network-centric approach, focusing on scheduling using ordered device proximity groups. This approach should help extend network lifespan as more devices are introduced by minimizing redundancy and distributing resource usage. The approach should also enhance network resilience, as devices can dynamically adjust their schedules if a neighbor fails. Given a sufficiently dense network, devices in proximity groups can activate each other based on forecasted event patterns/locations [4], increasing responsiveness for out-of-schedule and non-stationary events.

This effort forms part of an ongoing project to build a new low-power, affordable multi-sensor device for biodiversity monitoring, facilitating the large-scale gathering of rich data essential for maintaining conservation initiatives. .

⁶The SAFE project

References

- [1] Rachel T. Buxton, Patrick E. Lendrum, Kevin R. Crooks, and George Wittemyer. 2018. Pairing camera traps and acoustic recorders to monitor the ecological impact of human disturbance. *Global Ecology and Conservation* 16 (2018), e00493. <https://doi.org/10.1016/j.gecco.2018.e00493>
- [2] Vincent Droissart, Laura Azandi, Eric Rostand Onguene, Marie Savignac, Thomas B. Smith, and Vincent Deblauwe. 2021. PICT: A low-cost, modular, open-source camera trap system to study plant–insect interactions. *Methods in Ecology and Evolution* 12, 8 (2021), 1389–1396. <https://doi.org/10.1111/2041-210X.13618> arXiv:<https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.13618>
- [3] Alaa Hafez, Amgad Saaed, Mohamed Emam, Mohamed Hassan, and Mohamed Shalaby. 2023. An Optimized Radar Sensors Schedule Using Genetic Algorithms. In *2023 International Telecommunications Conference (ITC-Egypt)*. 79–82. <https://doi.org/10.1109/ITC-Egypt58155.2023.10206163>
- [4] Becky E Heath, Neel P Le Penru, James Skinner, C. David L. Orme, Sarab S Sethi, Robert M Ewers, and Lorenzo Picinali. 2024. MAARU: Multichannel Acoustic Autonomous Recording Unit for spatial ecosystem monitoring. *bioRxiv* (2024). <https://doi.org/10.1101/2024.01.23.576628> arXiv:<https://www.biorxiv.org/content/early/2024/01/26/2024.01.23.576628.full.pdf>
- [5] Andrew P. Hill, Peter Prince, Evelyn Piña Covarrubias, C. Patrick Doncaster, Jake L. Snaddon, and Alex Rogers. 2018. AudioMoth: Evaluation of a smart open acoustic device for monitoring biodiversity and the environment. *Methods in Ecology and Evolution* 9, 5 (2018), 1199–1211. <https://doi.org/10.1111/2041-210X.12955> arXiv:<https://besjournals.onlinelibrary.wiley.com/doi/pdf/10.1111/2041-210X.12955>
- [6] Seunghyeok Jeon, Yonghun Choi, Yeonwoo Cho, and Hojung Cha. 2023. HarvNet: Resource-Optimized Operation of Multi-Exit Deep Neural Networks on Energy Harvesting Devices. In *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services (Helsinki, Finland) (MobiSys '23)*. Association for Computing Machinery, New York, NY, USA, 42–55. <https://doi.org/10.1145/3581791.3596845>
- [7] Bo Jiang, Binoy Ravindran, and Hyeonjoong Cho. 2013. Probability-Based Prediction and Sleep Scheduling for Energy-Efficient Target Tracking in Sensor Networks. *IEEE Transactions on Mobile Computing* 12, 4 (2013), 735–747. <https://doi.org/10.1109/TMC.2012.44>
- [8] Stefan Kahl, Connor M. Wood, Maximilian Eibl, and Holger Klinck. 2021. BirdNET: A deep learning solution for avian diversity monitoring. *Ecological Informatics* 61 (2021), 101236. <https://doi.org/10.1016/j.ecoinf.2021.101236>
- [9] Laretta Andrew Laneng, Futoshi Nakamura, Yasuyuki Tachiki, and Charles S. Vairappan. 2021. Camera-trapping assessment of terrestrial mammals and birds in rehabilitated forest in INIKEA Project Area, Sabah, Malaysian Borneo. *Landscape and Ecological Engineering* 17, 2 (2021), 135–146. <https://doi.org/10.1007/s11355-020-00442-7>
- [10] James Long and Oral Buyukozturk. 2020. Collaborative duty cycling strategies in Energy Harvesting Sensor Networks. <https://dspace.mit.edu/handle/1721.1/126602>
- [11] Di Tian and Nicolas D. Georganas. 2002. A coverage-preserving node scheduling scheme for large wireless sensor networks. In *Proceedings of the 1st ACM International Workshop on Wireless Sensor Networks and Applications* (Atlanta, Georgia, USA) (WSNA '02). Association for Computing Machinery, New York, NY, USA, 32–41. <https://doi.org/10.1145/570738.570744>
- [12] Christopher J. C. H. Watkins and Peter Dayan. 1992. Q-learning. *Machine Learning* 8, 3 (01 May 1992), 279–292. <https://doi.org/10.1007/BF00992698>
- [13] Connor M. Wood and Stefan Kahl. 2024. Guidelines for appropriate use of BirdNET scores and other detector outputs. *Journal of Ornithology* 165, 3 (2024), 777–782. <https://doi.org/10.1007/s10336-024-02144-5>
- [14] J. Wolfgang Wägele and Georg F. Tschan. 2024. Weather stations for biodiversity: a comprehensive approach to an automated and modular monitoring system. *Advanced Books* 1 (2024). Advanced Books. <https://doi.org/10.3897/ab.e119534> arXiv:<https://doi.org/10.3897/ab.e119534>
- [15] Dirk Zeuss, Lisa Bald, Jannis Gottwald, Marcel Becker, Hicham Bellafkir, Jörg Bendix, Phillip Bengel, Larissa T. Beumer, Roland Brandl, Martin Brändle, Stephan Dahlke, Nina Farwig, Bernd Freisleben, Nicolas Friess, Lea Heidrich, Sven Heuer, Jonas Höchst, Hajo Holzmann, Patrick Lampe, Martin Leberecht, Kim Lindner, Juan F. Masello, Jonas Mielke Möglich, Markus Mühlhling, Thomas Müller, Alexey Noskov, Lars Oppenoorth, Carina Peter, Petra Quillfeldt, Sascha Rösner, Raphaël Royauté, Christian Mestre-Runge, Dana Schabo, Daniel Schneider, Bernhard Seeger, Elliot Shayle, Ralf Steinmetz, Pavel Tafo, Markus Vogelbacher, Stephan Wöllauer, Sohaib Younis, Julian Zobel, and Thomas Nauss. 2024. Nature 4.0: A networked sensor system for integrated biodiversity monitoring. *Global Change Biology* 30, 1 (2024), e17056. <https://doi.org/10.1111/gcb.17056> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/gcb.17056> e17056 GCB-23-1269.R1.
- [16] Filip Zápala and Miroslav Kasal. 2014. Using the Goertzel algorithm as a filter. In *2014 24th International Conference Radioelektronika*. 1–3. <https://doi.org/10.1109/Radioelek.2014.6828441>