# Experimental Evaluation of the Performance of UAV-assisted Data Collection for Wake-up Radio-enabled Wireless Networks

Abhimanyu V. Sheshashayee,* Matteo Bordin,* Pietro Brach del Prever,* Davide Villa,*
Hai Cheng,* Chiara Petrioli,† Tommaso Melodia,* Stefano Basagni*
*Institute for the Wireless Internet of Things at Northeastern University, Boston, MA, U.S.A.
E-mail: {sheshashayee.a, bordin.m, brachdelprever.p, villa.d, cheng.hai, melodia, s.basagni}@northeastern.edu
†Department of Computer, Control and Management Engineering, Università di Roma "La Sapienza," Rome, Italy
E-mail: petrioli@diag.uniroma1.it

***Abstract*—Wireless Sensor Networks (WSNs) are pivotal in various applications, including precision agriculture, ecological surveillance, and the Internet of Things (IoT). However, energy limitations of battery-powered nodes are a critical challenge, necessitating optimization of energy efficiency for maximal network lifetime. Existing strategies like duty cycling and Wake-up Radio (WuR) technology have been employed to mitigate energy consumption and latency, but they present challenges in scenarios with sparse deployments and short communication ranges. This paper introduces and evaluates the performance of Unmanned Aerial Vehicle (UAV)-assisted mobile data collection for WuR-enabled WSNs through physical and simulated experiments. We propose two one-hop UAV-based data collection strategies: a naïve strategy, which follows a predetermined fixed path, and an adaptive strategy, which optimizes the collection route based on recorded metadata. Our evaluation includes multiple experiment categories, measuring collection reliability, collection cycle duration, successful data collection time (latency), and node awake time to infer network lifetime. Results indicate that the adaptive strategy outperforms the naïve strategy across all metrics. Furthermore, WuR-based scenarios demonstrate lower latency and considerably lower node awake time compared to duty cycle-based scenarios, leading to several orders of magnitude longer network lifetime. Remarkably, our results suggest that the use of WuR technology alone achieves unprecedented network lifetimes, regardless of whether data collection paths are optimized. This underscores the significance of WuR as the technology of choice for all energy critical WSN applications.***

## I. Introduction

Networks of wireless devices acting as interface to the physical world, known as *Wireless Sensor Networks* (WSNs), have become the communication backbone of countless applications. They have been applied to various domains, including precision agriculture, ecological surveillance, and the *Internet of Things* (IoT) [1]. The crux of wireless sensor networking resides in its scale, its pervasiveness, and in being able to communicate wirelessly. This is because WSN devices (or *nodes*) can be deployed virtually anywhere, independently of the presence of infrastructure (e.g., wired networks) or power. However, being battery-powered poses limitations on energy availability, which consequently limits a node's lifetime and the overall duration of network operations (*network lifetime*). In most scenarios, the replacement of nodes' batteries is inconvenient or entirely impossible. Thus, an important problem for WSNs is that of optimizing energy efficiency in order to attain maximal network lifetime.

Approaches for optimizing energy efficiency have been explored at all layers of a node architecture [2]. In contemporary designs, the energy consumed by communication is found to be significantly higher than by all other operations [3]. The majority of energy consumption by communication is due to the radio's *idle listening*. This wasteful energy consumption can be reduced by the use of *duty cycling* [4], wherein the node's radio is kept off (*asleep*), and periodically turned on (*awake*) for brief intervals, during which communication occurs. In the asleep state, the energy consumption is orders of magnitudes lower than in the awake state (usually microwatts vs. milliwatts). While this approach can saliently reduce energy consumption due to idle listening, it noticeably increases the expected data delivery times (end-to-end latency) [5]. These high latencies can be ameliorated by making the transition to the awake state occur "on demand" using *Wake-up Radio* (WuR) technology [6], [7], [8]. In addition to its main radio, each node is endowed with an ultra-low-power auxiliary radio that remains always on. When communication with a particular node is required, that node can be awoken by sending it a *Wake-up Sequence* (WuS) that matches its *Wake-up Address* (WuA). After exchanging packets, nodes return to sleep. This approach significantly attenuates both energy consumption and latency [9].

Perhaps the most important function of a WSN is that of *data collection*. Nodes obtain sensor readings from the environment and collate them into data packets that need to be transported to a collection point (the network *sink*). In many WSN scenarios data is collected by routing packets through multiple nodes (multi-hop routing) [10]. However, routing-based data collection requires the network to be fully connected, which prevents applicability to scenarios with sparse deployments [11]. This is further exacerbated by the use of WuR technology, as contemporary WuR designs typically have relatively short communication ranges ($\leq$ 25 m) [12], [13].

Since a WuR-enabled WSN would need to be fully connected with respect to WuR links in order to route data packets, this imposes an even greater constraint on network density, which affects deployment and operational costs.

Another approach to data collection envisions sending a *mobile data collector* to the node, to collect data directly from the source via a simple one-hop wireless transmission [14], [15], [16]. The collector can be any kind of mobile device, outfitted with communication equipment and an on-board controller. Sufficient mobility, like that of an *Unmanned Aerial Vehicle* (UAV) [17], significantly reduces the difficulty of collecting data from remote nodes [18]. A further benefit of mobile data collection is the avoidance of energy consumption by routing. Even if a route from a node to the sink can be found, multi-hop routing always consumes more energy than one-hop forwarding [19], [20]. This leads to shorter node lifetime, and in turn to shorter network lifetime.

In this paper, we evaluate the performance of UAV-assisted mobile data collection for WuR-enabled WSNs via physical and simulated experiments. Our aim is to demonstrate that the joint exploitation of WuR technology and one-hop forwarding produces network lifetimes that are unthinkable in duty cycling-based WSNs. We start by designing two one-hop UAV-based data collection strategies. The first, a *naïve* strategy, keeps no information about node locations and links (namely, on the network topology), and visits the nodes according to a predetermined fixed path. We then define an *adaptive* strategy that aims at optimizing the collection route based on metadata recorded during previous collection cycles. Our evaluation is based on multiple categories of experiments in which the performance of UAV-based data collection is evaluated. We use WuR-enabled wireless sensor nodes and a quad-rotor drone UAV to conduct physical experiments that evaluate both strategies in networks with duty cycling and with WuR. Parameters and results from these physical experiments are used to inform and validate simulations, which we use to examine the effects of scale on performance. At each collection cycle, we measure the percentage of nodes from which data packets are collected (reliability), the duration of the collection cycle, and the time of successful data collection (latency). We also measure the node awake time, which we use to infer network lifetime.

Results show that the adaptive strategy obtains the best performance for all metrics. The time spent awake by the nodes is considerably lower for the WuR-based scenarios than for duty cycle-based networks. The latency of WuR-based scenarios is generally lower than that in scenarios with shorter duty cycles. Overall, the lifetime of networks with WuR technology, whether data collection paths are optimized or not, is shown to be several orders of magnitude longer than that of networks with duty cycle. This makes WuRs the technology of choice for all energy-critical WSN applications.

The paper is organized as follows. Section II describes the experimental scenarios and the data collection strategies. Section III lists performance metrics, experimental parameters and discusses results. Section IV concludes the paper.

## II. UAV-assisted Data Collection

### A. Network Scenarios

We consider scenarios that consist of $N$ wireless nodes scattered in an $L \times W$ deployment region. Each node is endowed with a wireless transceiver (the *main radio*), used for data communications. A UAV with a low-power MagoNode++ [21], is used to collect data from the network nodes. It launches from a *Base Station* (BS), traverses the region in which the nodes are deployed, stops at designated *collection points* and wirelessly requests and gathers data from every wireless node in the deployment region, finally returning to the BS to recharge/refuel. This process is known as a *collection cycle*. A powered sink is situated in or near the deployment region. The sink sports a relatively long-range radio to communicate with the UAV. As the UAV collects data packets, it relays them to the sink. Upon returning to the BS, the UAV batteries are checked and replaced if needed.

We consider two types of scenarios:

- In *duty cycle-based* scenarios each node operates its main radio according to a preset duty cycle $d$. Following a duty cycle of $d\%$ means that a node radio remains awake for $d/100$ s and goes to sleep for $(1 - d/100)$ s, repeating indefinitely. When the UAV hovers by a collection point for a time $\tau$, it unicasts a Request-to-Receive (RTR) control packet to the node (via its main radio). If the node is awake, it responds with a data packet.
- In *WuR-based* scenarios each node also feature a WuR, and is set to wake up according to a unique ID [22], [23]. A node keeps its main radio asleep until it receives a wake-up sequence (via its WuR) matching its ID. When awoken, the node sends a data packet to the UAV using the main radio, and then returns to sleep.

In both scenarios the UAV *unicasts* its requests (whether an RTR packet or a wake-up sequence), which requires the UAV to know the unique address of the node to whom the request is sent. This might be known at network deployment, or can become known after a few connection cycles (the latter is the case of our experiments below). Using unicast requests makes the data collection more robust and simplifies channel access.

### B. Data Collection Strategies

We consider two collection strategies: *naïve* and *adaptive*.

- The naïve strategy considers the node locations to be unknown. The UAV is set to fly in a rectangular spiral path that covers the entire deployment region.[1] The UAV periodically stops at preset collection points along this path. At each collection point, the UAV attempts to collect data from all nodes within its transmission reach.
- The adaptive strategy (AS) records topology metadata, which are used to estimate node locations, shorten the

[1] Traveling on rectangular spiral paths has been shown to produce better performance results than following other naïve paths [24], [25].

flight path for subsequent collection cycles, and eventually converge to a near-optimal flight path.[2] The UAV initially starts with a naïve collection cycle, but keeps a record of which nodes were successfully collected from at each collection point. This information is used to estimate node locations, using supervised likelihood-maximization assisted by Nelder-Mead search. Based on the estimated node locations, a sequence of $N$ new collection points is determined (one collection point for each node) using a pattern-search adaptation of the Nearest Neighbor Algorithm [27]. The strategy converges to the near-optimal path when the mean Euclidean distance between current and previous estimates of node locations is below a preset threshold. Once the strategy has converged to the near-optimal path, it will stop making adjustments to subsequent paths. (All future collection cycles will take the same path).

Fig. 1 illustrates both strategies on a sample topology with 4 nodes deployed in a square area. Fig. 1a depicts the topology: The BS location is depicted at the center of the area, represented by a small pentagon; the nodes are depicted as small gray rectangles. Fig. 1b presents the path taken and the preset collection points (small green dots) visited by the UAV during the naïve collection cycle. In the case of the adaptive collection strategy (Figures 1c to 1f), the UAV initially performs the same naïve collection cycle. However, Fig. 1c shows the metadata recorded during this initial collection cycle (small colored dots). Fig. 1d shows the estimated node locations following the initial collection cycle, as indicated by the center of each circle surrounding a cluster of corresponding collection points. Fig. 1e presents the new sequence of collection points and the path taken by the UAV for the second adaptive collection cycle. Finally, Fig. 1f shows the collection points of the path to which the adaptive strategy converges.

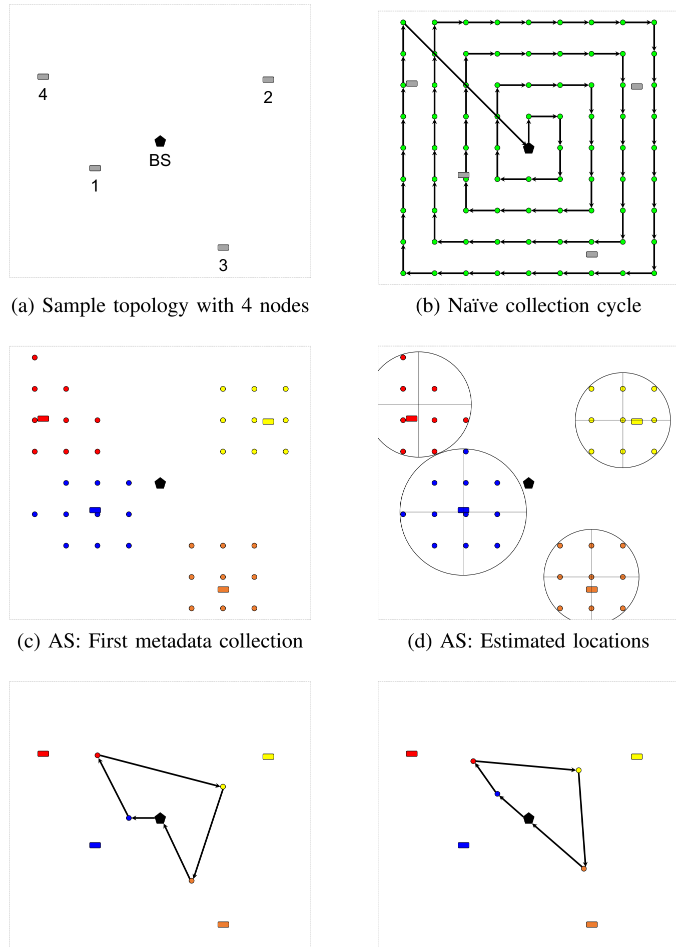

(a) Sample topology with 4 nodes (b) Naïve collection cycle
(c) AS: First metadata collection (d) AS: Estimated locations
(e) AS: First adapted path (f) AS: Near-optimal path

Fig. 1. (a) A sample topology. (b) The naïve collection strategy. (c)–(f) The adaptive collection strategy (AS).

## III. Experimental Evaluation

Our experiments have the overarching goal of showing the effectiveness of using WuR technology for UAV-assisted data collection. We conducted two types of experiments, namely, *testbed-based experiments*, where we flew a UAV to collect data from ten wireless nodes (Section III-A), and *simulation-based experiments*, which we validated via the testbed-based results, and which we used to investigate the effect of scale on UAV-assisted data collection (Section III-B).

We investigate the following metrics (all averages).

- *Packet Delivery Ratio (PDR)*: The percentage of all generated data packets that are successfully delivered to the sink via the UAV.
- *Awake time*: The time (in seconds) for which a node remains awake during a collection cycle, averaged over all nodes in the network.[3]
- *Latency*: The time (in seconds) taken to successfully collect a data packet and forward it to the sink. This metric includes the time for waking up nodes using the WuR, data transmission from nodes to UAV, and data forwarding from the UAV to the sink.
- *Network lifetime*: The time (in days) from the start of network operations until the first device runs out of energy.[4] (We stipulate that devices are continuously visited by a UAV: Once one UAV has finished a visit to the device, another one follows for a new collection cycle.)

Common parameter settings for both types of experiments include the following. Data are produced by the nodes according to a Poisson distribution with mean 1 minute. Each data packet is 70 B long, accounting for the application payload and for the headers added by lower layers (simple MAC and physical layer). The UAV flies at a speed of 2.15 m/s, and

[2] Determining optimal paths would require solving the NP-hard Travelling Salesman Problem. Near-optimal paths are determined by using a constructive heuristic that approximates the optimal path in sub-quadratic time [26].

[3] As the energy consumption of a node is dominated by idle listening, the awake time gives an indication of the energy consumed over the course of our experiments. Measuring actual energy values proved problematic in testbed-based experiments.

[4] This definition is consistent with that of prior work about the lifetime of wireless networks. Although the network may still serve its intended purpose even if some of its devices are "dead," this conservative definition provides an informative lower bound on the network lifetime [28].

stops at each collection point for a time $\tau = 4$ s. In duty cycle-based (WuR-based) scenarios, the RTR packet (wake-up sequence) is transmitted for a total of 25 (20) times, equally spaced within one second of the 4 that the UAV resides at the collection point. The UAV is capable of self-localization via high-precision onboard GPS [17]. In all of our experiments, flight altitude is set to 5 m, which ensures robust and reliable WuR-based communication [13]. The wireless nodes used in our testbed-based experiments are the WuR-endowed ultra-low-power *MagoNode++*, that can be used in both duty cycle and WuR mode [21]. The MagoNode++ has been modeled in detail for our simulation experiments [7], [9]. Accordingly, the data rate of the main radio is 250 kbps, and the transmission range is 70 m. In duty cycled-based scenarios, the duty cycle $d$ is varied in the set $\{100, 50, 10, 5\}\%$. Once the UAV is at a collection point it broadcasts a 3 B long RTR packet. In WuR-based scenarios the wake-up sequence is 8 b long. These sequences are transmitted at a data rate of 1 kbps. The transmission range of the WuR is set to 25 m [13].

### A. *Testbed-based Experiments*

As mentioned, the wireless nodes used in our testbed-based experiments are the WuR-endowed *MagoNode++* [21]. The UAV is a *Monarch* quad-rotor vehicle, outfitted with a *Pixhawk Mini* flight controller operating with *Ardupilot* firmware, and an Intel NUC *NUC7i7DN* for the on-board computer [17]. This model is fairly heavy, at approximately 5 kg, but offers an impressive flight time of over 40 minutes. The NUC is remotely accessed via a WiFi connection in order to run the program that controls flight and communication.

Our experiments were conducted at Northeastern University's *Expeditionary Cyber and Unmanned Aerial System* (ECUAS) lab at the Innovation Campus in Burlington, MA. Specifically, all experiments were conducted inside the outdoors netted enclosure. This enclosure is $45.72\text{ m} \times 60.96\text{ m} \times 18.29$ m in size. Ten wireless nodes were mounted on tripods and distributed across the ground. The BS was situated at ground-level in the center of the enclosure. The sink was situated inside the enclosure, plugged into the local power grid, listening for data packets from the UAV.

For the naïve collection strategy, we performed multiple independent collection cycles for both duty cycle-based scenarios and WuR-based scenarios. Experiments with the adaptive collection strategy were performed as multiple sets of collection cycles, starting with a naïve collection cycle and progressing until the strategy converged on a near-optimal path, which was then used to collect measurements.

▷ **Results**. The PDR is consistently 100%, given the relatively low data traffic generation rate and collision-free transmission thanks to the use of unicast for collecting data from the nodes. Table I lists the results of all other metrics for both types of scenarios averaged over all collection cycles.

Independently of the collection strategy, the awake time for the duty cycle-based scenarios decreases with the value $d$ of the duty cycle, as expected. Also expected, the awake time for the WuR-based scenario is substantially lower than that of any of the duty cycle-based scenarios: over 98% lower than in the case with $d = 5\%$, independently of the collection strategy. The adaptive strategy reduces the time taken per collection cycle. As a consequence, the number of collection cycles performed by the adaptive strategy is considerably greater than the number of collection cycles performed by the naive strategy. The reason why the awake times vary in networks with duty cycle is because we look at the amount of time spent awake per collection cycle. For example, for a duty cycle of 100%, the time spent awake is the same for both strategies (i.e., 100%). However, because the adaptive strategy results in a faster collection cycle, the number of seconds awake is lower for the adaptive strategy than for the naive strategy.

Latency in duty cycle-based scenarios clearly increases for decreasing values of $d$. We observe that latency in the WuR-based scenarios is higher than that in scenarios with 50% and 100% duty cycle. This is because of the WuR range, which is shorter than that of the main radio, thus allowing the UAV to awake and receive data only from nodes closer to it. When a node is further away from the UAV, the likelihood of successful awakening decreases [13]. For the remaining two values of the duty cycle that we considered, using a WuR is faster in retrieving data, because the longer sleep times imposed when $d = 5, 10\%$ makes the UAV wait longer.

Network lifetimes in WuR-based scenarios are always remarkably longer than for any duty cycle scenarios. [5] For the naïve strategy, the WuR-based scenario achieves over 1563 (78) times longer lifetime than in scenarios with $d = 100\%$ ($d = 5\%$). For the adaptive strategy, these ratios are over 2340 times and 117 times, respectively.

Using a collection strategy that optimizes the UAV path produces better performance of all measured metrics, irrespective of the scenario. This is expected, as being able to select collection points that are closer to the nodes imposes lower energy consumption and shortens latency. We notice, however, that the WuR-induced improvements of critical metrics such as network lifetime over duty cycling are so overwhelmingly high even when using naïve strategies that implementing complex path optimization techniques appears unnecessary.

### B. *Simulation-based experiments*

Simulations are conducted using the *GreenCastalia* simulator [29], an extension of the *Castalia* simulator [30] that is based on OMNeT++ [31]. The MagoNode++, the Monarch UAV and its mobility are implemented in GreenCastalia with parameters obtained from the testbed-based experiments.

Our simulation models and settings have been validated through testbed-based experiments. We implemented the testbed scenario described above in GreenCastalia. The simulation results on all metrics of interest are in agreement with

[5] As it is unfeasible to run testbed-based experiments for days at a time, the network lifetime is calculated using the energy model from the simulation-based experiments. Each node is powered by one AA lithium-ion battery with a capacity of 10656 J. The lifetime of each node is calculated as battery capacity divided by mean energy consumption during a collection cycle, multiplied by the mean duration of a collection cycle. The lowest node lifetime sets the lifetime of the network.

TABLE I
PERFORMANCE RESULTS FROM TESTBED-BASED EXPERIMENTS.

| **Scenarios** | NAÏVE COLLECTION STRATEGY | | | ADAPTIVE COLLECTION STRATEGY | | |
|---|---|---|---|---|---|---|
| | **Awake time [s]** | **Latency [s]** | **Network lifetime [days]** | **Awake time [s]** | **Latency [s]** | **Network lifetime [days]** |
| 100% duty cycle | $68.983 \pm 0.428$ | $0.073 \pm 0.008$ | 3.671 | $23.258 \pm 0.144$ | $0.050 \pm 0.005$ | 3.671 |
| 50% duty cycle | $33.478 \pm 0.275$ | $0.343 \pm 0.037$ | 7.341 | $12.198 \pm 0.100$ | $0.107 \pm 0.012$ | 7.272 |
| 10% duty cycle | $7.007 \pm 0.078$ | $0.538 \pm 0.069$ | 36.706 | $2.526 \pm 0.028$ | $0.135 \pm 0.017$ | 36.712 |
| 5% duty cycle | $4.185 \pm 0.030$ | $0.618 \pm 0.106$ | 73.413 | $2.441 \pm 0.018$ | $0.143 \pm 0.024$ | 73.401 |
| WuR | $0.076 \pm 0.007$ | $0.427 \pm 0.032$ | 5740.455 | $0.022 \pm 0.002$ | $0.125 \pm 0.009$ | 8592.477 |

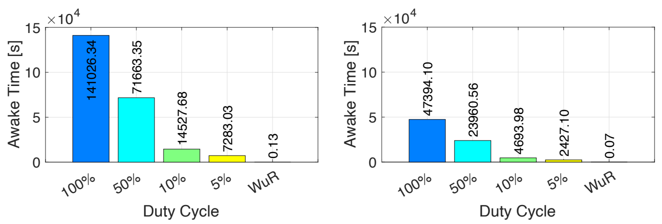


(a) Naïve strategy (b) Adaptive strategy

Fig. 2. Simulation-based experiments: Awake time in networks with $N = 128$ and $L, W = 500$ m.

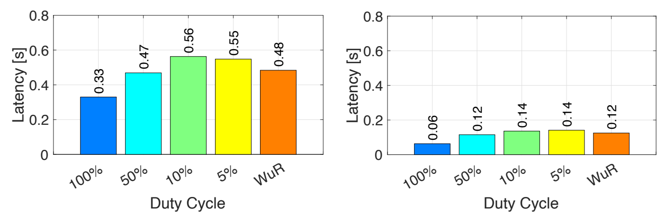


(a) Naïve strategy (b) Adaptive strategy

Fig. 3. Simulation-based experiments: Latency in networks with $N = 128$ and $L, W = 500$ m.

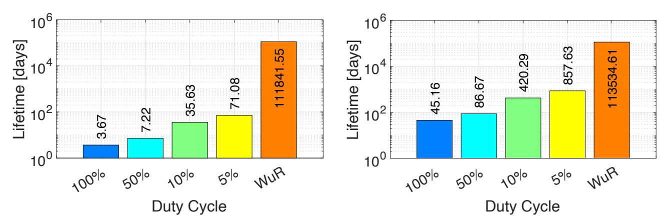


(a) Naïve strategy (b) Adaptive strategy

Fig. 4. Simulation-based experiments: Network lifetime in networks with $N = 128$ and $L, W = 500$ m.

those from the testbed-based experiments: All averages from simulations are within 5% of those from the testbed. All results shown here have been obtained by averaging the outcomes of a number of simulation runs sufficient to obtain 95% confidence with 5% precision.

In our simulations we varied the size of the network $N \in \{48, 64, 128\}$ and the size of the deployment area $L \times W \in \{(100 \text{ m} \times 100 \text{ m}), (250 \text{ m} \times 250 \text{ m}), (500 \text{ m} \times 500 \text{ m})\}$. As result trends are the same for all combinations obtained by varying these parameters, in the following we show results for networks with 128 nodes scattered in the largest deployment area, which is the case more representative of large scale.

▷ **Results**. The PDR of all collection cycles remains consistently 100%. The naïve strategy comprehensively covers the entire deployment region, making packet losses unlikely. If the adaptive strategy encounters a missed collection, it retries previously successful collection points until it can recover the packet, which again results in low likelihood of packet loss.

Fig. 2 shows results concerning the awake time for the naïve (Fig. 2a) and the adaptive (Fig. 2b) collection strategies. Results for the awake time mimic those observed in the testbed-based experiments, in that this metric decreases with the values of the duty cycle $d$ in duty cycle-based scenarios, and is remarkably better in WuR-based scenarios. For instance, nodes with a WuR stay awake 99% less than when they operate according to a 5% duty cycle, irrespective of the collection strategy. This suggests that scale has minimal impact on the awake time of wireless nodes with or without WuR technology.

Fig. 3 shows results concerning the latency for the naïve (Fig. 3a) and the adaptive (Fig. 3b) collection strategies. As observed in the case of testbed-based experiments, latency in WuR-enabled networks improves only compared to latency in networks with low duty cycles. This, again, is a consequence of the lower range of WuR technology with respect to that of the main radio (25 m vs. 70 m), allowing networks with duty cycles to reach directly more nodes.

Fig. 4 shows results concerning the network lifetime for the naïve (Fig. 4a) and the adaptive (Fig. 4b) collection strategies. In duty cycle-based scenarios the network lifetime grows noticeably with decreasing values of $d$, irrespective of the collection strategy. This is clearly because nodes stay awake less and less, and therefore consume less, thus lasting longer. The improvement in network lifetime obtained by using WuR technology is remarkable. For instance, with respect to networks that duty cycle with $d = 5\%$, networks with WuRs last 1575 (132) times more if the UAV follows a naïve (adaptive) path.

Again, the adaptive strategy obtains improvements in all considered metrics, particularly *in networks with duty cycling*. Improvements of the adaptive strategy over the naïve one are less remarkable in networks with WuRs, due to the impressive energy efficiency gains obtained by using WuR technology, for which nodes remain awake only as long as absolutely necessary for communicating. This is particularly evident for the network lifetime, which sees an improvement of at least one order of magnitude in networks with duty cycling, but below a mere $2\%$ in WuR-based networks. This, again, suggests that the energy savings afforded by using WuRs are so remarkable to be independent of the particular data collection strategy implemented by the UAV, favoring the use of simple and cost-effective data collection strategies.

## IV. Conclusions

In this study, we explore the efficiency of UAV-assisted data collection in WSNs using duty cycling or WuR technology for energy conservation. We introduced a basic collection strategy and developed an adaptive approach for UAV data gathering. Through experimental and simulation-based evaluations in realistic network scenarios, we found that WuR technology significantly outperforms duty cycling in all key metrics. Notably, the energy savings provided by WuRs are substantial enough to potentially render complex UAV route optimization strategies unnecessary.

## Acknowledgments

Chiara Petrioli was supported in part by the Italian MUR-funded PON project SMARTOUR.

## References

[1] B. Nayak, S. K. Pani, T. Choudhury, S. Satpathy, and S. N. Mohanty, Eds., *Wireless Sensor Networks and the Internet of Things: Future Directions and Applications*. Apple Academic Press, September 2021.

[2] G. Popli and R. P. Singh, "Optimization of energy consumption in wireless sensor network – a review," in *Proceedings of GUCON 2018*, Greater Noida, Uttar Pradesh, India, September 28–29 2018, pp. 961–967.

[3] S. Basagni, V. Di Valerio, G. Koutsandria, and C. Petrioli, "On the impact of local computation over routing performance in green wireless networks," in *Proceedings of IEEE WoWMoM 2018*, Chania, Greece, October 12–15 2018, pp. 1–9.

[4] R. C. Carrano, D. Passos, L. C. S. Magalhaes, and C. V. N. Albuquerque, "Survey and taxonomy of duty cycling mechanisms in wireless sensor networks," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, pp. 181–194, June 2014.

[5] S. Basagni, F. Ceccarelli, F. Gattuso, and C. Petrioli, "Demo abstract: Abating LPL-induced latency with wake-up radio technology," in *Proceedings of ACM IoT DI 2017*, Pittsburg, PA, April 18–21 2017, pp. 1–2.

[6] C. Petrioli, D. Spenza, P. Tommasino, and A. Trifiletti, "A novel wake-up receiver with addressing capability for wireless sensor nodes," in *Proceedings of IEEE DCOSS 2014*, Marina Del Rey, CA, May 26–28 2014, pp. 18–25.

[7] D. Spenza, M. Magno, S. Basagni, L. Benini, M. Paoli, and C. Petrioli, "Beyond duty cycling: Wake-up radio with selective awakenings for long-lived wireless sensing systems," in *Proceedings of IEEE INFOCOM 2015*, Hong Kong, China, April 26–30 2015, pp. 522–530.

[8] R. Piyare, A. L. Murphy, C. Kiraly, P. Tosato, and D. Brunelli, "Ultra low power wake-up radios: A hardware and networking survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2117–2157, July 24 2017.

[9] A. V. Sheshashayee, C. Petrioli, and S. Basagni, "On the effectiveness of semantic addressing for wake-up radio-enabled wireless sensor networks," in *Proceedings of IEEE PIMRC 2022*, Virtual, September 12–15 2022, pp. 1–6.

[10] C. Nakas, D. Kandris, and G. Visvardis, "Energy efficient routing in wireless sensor networks: A comprehensive survey," *Algorithms*, vol. 13, no. 3, pp. 1–65, 2020.

[11] C. Luo, M. N. Satpute, D. Li, Y. Wang, W. Chen, and W. Wu, "Fine-grained trajectory optimization of multiple UAVs for efficient data gathering from WSNs," *IEEE/ACM Transactions on Networking*, vol. 29, no. 1, pp. 162–175, October 9 2021.

[12] S. Basagni, F. Ceccarelli, C. Petrioli, N. Raman, and A. V. Sheshashayee, "Wake-up radio ranges: A performance study," in *Proceedings of IEEE WCNC 2019*, Marrakech, Morocco, April 15–19 2019, pp. 1–5.

[13] A. V. Sheshashayee, J. Buczek, C. Petrioli, and S. Basagni, "Experimental evaluation of wake-up radio ranges for UAV-assisted mobile data collection," in *Proceedings of IEEE WCNC 2022*, Austin, TX, U.S.A., April 10–13 2022, pp. 716–721.

[14] S. Basagni, A. Carosi, C. Petrioli, and C. A. Phillips, "Coordinated and controlled mobility of multiple sinks for maximizing the lifetime of wireless sensor networks," *ACM/Springer Wireless Networks*, vol. 17, no. 3, pp. 759–778, April 2011.

[15] Y. Gu, F. Ren, Y. Ji, and J. Li, "The evolution of sink mobility management in wireless sensor networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 507–524, First quarter 2016.

[16] C. Zhan, Y. Zeng, and R. Zhang, "Energy-efficient data collection in UAV enabled wireless sensor network," *IEEE Journal on Selected Areas in Communications*, vol. 7, no. 3, pp. 328–331, June 2018.

[17] J. Buczek, L. Bertizzolo, S. Basagni, and T. Melodia, "What is a wireless UAV? A design blueprint for 6G flying wireless nodes," in *Proceedings of ACM WiNTECH 2021*, New Orleans, LA, January 31–February 4 2022, pp. 1–7.

[18] Z. Wei, M. Zhu, N. Zhang, L. Wang, Y. Zou, Z. Meng, H. Wu, and Z. Feng, "UAV assisted data collection for Internet of Things: A survey," *IEEE Internet of Things Journal*, vol. 9, no. 15, pp. 15 460–15 483, May 2022.

[19] S. Smruthi, R. S. Krishna, and M. Panda, "Low energy sensor data collection using unmanned aerial vehicles," in *Proceedings of ICOEI 2019*, Tirunelveli, Tamil Nadu, India, June 3–5 2019, pp. 740–745.

[20] S. Basagni, G. Koutsandria, and C. Petrioli, "Enabling the mobile IoT: Wake-up unmanned aerial systems for long-lived data collection," in *Proceedings of IEEE MASS 2019*, Monterey, CA, November 4–7 2019, pp. 1–8.

[21] M. Paoli, D. Spenza, C. Petrioli, M. Magno, and L. Benini, "MagoNode++: A wake-up radio-enabled wireless sensor mote for energy-neutral applications," in *Proceedings of ACM/IEEE IPSN 2016*, Vienna, Austria, April 11–14 2016, pp. 1–2.

[22] A. V. Sheshashayee and S. Basagni, "WiEnum: Node enumeration in wireless networks," in *Proceedings of MILCOM 2015*, Tampa, FL, U.S.A., October 26–28 2015, pp. 1027–1032.

[23] A. V. Sheshashayee and S. Basagni, "WiLE: Leader election in wireless networks," *Ad Hoc & Sensor Wireless Networks*, vol. 44, no. 1–2, pp. 59–81, June 2019.

[24] V.-A. Vladuta, A.-M. Ghimes, C. Grumazescu, and V.-V. Patriciu, "Trajectory influence on data collection in wireless sensor networks using an unmanned aerial vehicle," in *Proceedings of ECAI 2018*, Iasi, Romania, June 28–30 2018, pp. 1–6.

[25] J. R. Antunes, L. Brisolara, and P. R. Ferreira, "UAVs as data collectors in the WSNs: Investigating the effects of back-and-forth and spiral coverage paths in the network lifetime," in *Proceedings of SBESC 2020*, Florianopolis, Brazil, November 23–27 2020, pp. 1–8.

[26] A. V. Sheshashayee, "Wake-up radio-enabled wireless networking: Measurements and evaluation of data collection techniques in static and mobile scenarios," PhD thesis, Northeastern University, Boston, MA, July 21 2022.

[27] I. Benmad, E. Driouch, and M. Kardouchi, "Data collection in UAV-assisted wireless sensor networks powered by harvested energy," in *Proceedings of IEEE PIMRC 2021*, Helsinki, Finland, September 13–16 2021, pp. 1351–1356.

[28] I. Dietrich and F. Dressler, "On the lifetime of wireless sensor networks," *ACM Transactions of Sensor Networks*, vol. 5, no. 1, pp. 1–38, January 2009.

[29] D. Benedetti, C. Petrioli, and D. Spenza, "GreenCastalia: An energy-harvesting-enabled framework for the Castalia simulator," in *Proceedings of ACM ENSSys 2013*, Rome, Italy, November 11–14 2013, pp. 1–6.

[30] A. Boulis, "Castalia: Revealing pitfalls in designing distributed algorithms in WSN," in *Proceedings of ACM SenSys 2007*, Sydney, Australia, November 6–9 2007, pp. 407–408.

[31] A. Varga, "Using the OMNeT++ discrete event simulation system in education," *IEEE Transactions on Education*, vol. 42, no. 4, p. 11 pp., November 1999.