Observability-Guided Node Scheduling for Long-Term Reconstruction Service in Energy-Constrained Wireless Sensor Networks

Guoqing Lu
South East Technological University
Waterford, Ireland
Email: guoqing1096@outlook.com

Jiaqi Zhou
South East Technological University
Waterford, Ireland
Email: 1794671582abc@gmail.com

***Abstract*—Long-lived wireless sensor networks (WSNs) may remain alive while losing the sensing directions needed for field reconstruction. This paper studies node scheduling as a reconstruction-service problem rather than a hardware-survival problem. The monitored field is represented by a finite-dimensional discrete-time linear dynamical system, and service quality is measured by the minimum eigenvalue of a finite-horizon observability Gramian. This observability margin is used as a service boundary for deciding whether the scheduled active set remains structurally informative for reconstruction.**

The proposed scheduling framework constructs a greedy minimum-feasible active set and rotates sensing responsibility according to residual energy, so that structurally informative sensors are not repeatedly exhausted. Its contribution is the reconstruction-service objective and its observability-feasibility-based integration with long-term energy rotation, with Gramian evaluation and greedy selection used as established primitives. Lossy delivery is treated as a scope check by evaluating delivered observations separately from scheduled observability; no post-delivery observability guarantee is claimed. In a representative 80-node long-term setting, the method reduces the horizon-averaged active ratio from 0.419 to 0.264, lowers RMSE from 0.1325 to 0.1005, and extends T_{obs} from 67.0 to 158.5 steps. Trace-driven and lossy-delivery checks further show that energy lifetime and packet delivery remain incomplete objectives when reconstruction service is the quantity to be preserved.

***Index Terms*—Wireless sensor networks; Observability; Node scheduling; State estimation; Energy efficiency**

I. Introduction

Wireless sensor networks (WSNs) combine distributed sensing with tight energy constraints [1], while state-estimation tasks often require selecting informative subsets of sensor measurements [2]. In this setting, node scheduling is a decision about information geometry: each activation changes the sensing subspace available to the estimator.

Conventional lifetime metrics do not fully capture this service requirement. A node being alive does not imply that useful information is available at the sink: coverage, connectivity, and residual-energy balance may prolong operation without preserving the sensing directions needed for state recovery. A network can therefore remain operational in hardware terms while its measurements collapse onto weakly informative directions. Lossy wireless delivery further sharpens this distinction because a scheduled measurement may not reach the estimator. For reconstruction-oriented WSNs, usefulness depends not only on delivery but also on whether the selected measurements preserve the information geometry of the latent state.

This paper formalizes reconstruction-service lifetime through finite-horizon observability. Observability-Gramian criteria are established tools for sensor selection [3], while the minimum eigenvalue of the Gramian has also been used as an observability-assurance measure [4]. We therefore quantify reconstruction service by the minimum eigenvalue of a finite-horizon observability Gramian and use it as a long-horizon service boundary under battery depletion, integrating that boundary with battery-aware load redistribution and established greedy-selection primitives.

The resulting framework activates only the sensors needed to satisfy an observability threshold and rotates that responsibility using residual energy. The core mechanism protects informative sensing directions from early depletion. Link quality is introduced separately to evaluate delivery efficiency. The main claim remains tied to scheduled-set observability rather than to a guarantee about the random post-delivery set; communication loss is therefore a boundary condition, not the primary optimization target.

The contributions are as follows:

- It formulates reconstruction-service lifetime as a WSN scheduling objective, separating node survival from sustained state reconstructability.
- It develops a lightweight observability-feasibility scheduling layer that constructs a greedy minimum-feasible active set and rotates sensing responsibility according to residual energy.
- It evaluates this system-level integration against representative energy, coverage, and estimation baselines, with lossy delivery used as a scope check rather than a post-delivery guarantee.

II. Related Work

WSN lifetime optimization is commonly formulated through node lifetime, coverage/connectivity, transmission, or

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

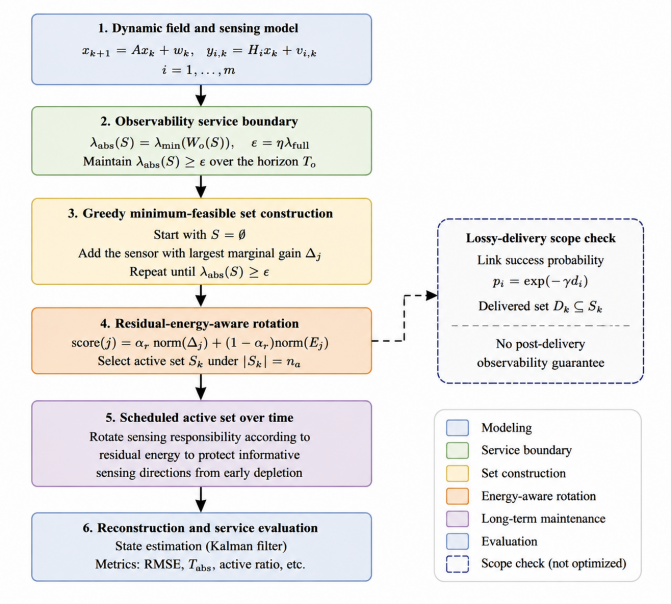


Fig. 1: Overview of the observability-guided scheduling framework. A greedy minimum-feasible active set is constructed from the observability boundary and then rotated according to residual energy to preserve long-term reconstruction service. Lossy delivery is evaluated separately as a scope check.

broader resource-allocation objectives [5]. These criteria are appropriate for maintaining network operation, but they do not directly guarantee that the surviving active set retains the sensing directions required for state reconstruction.

Energy-constrained WSN sensor selection has been formulated around Kalman-estimation performance [6]. Randomized greedy methods have also been developed for resource-constrained sensor selection under weak-submodular objectives [7], while online sparse sensor scheduling with explicit performance guarantees has been studied for time-varying systems [8]. These formulations establish important sensor-selection primitives, but our focus is repeated observability-feasible scheduling under battery depletion, where sensing responsibility must also be redistributed over time.

Update-scheduling work has additionally valued measurements according to their timeliness relative to downstream decision points [9]. In reconstruction-oriented WSNs, however, delivery alone is insufficient because the information value of a received measurement also depends on sensing geometry. We therefore treat lossy delivery as a separate efficiency boundary rather than as the primary service guarantee. The remaining gap is the joint treatment of observability, depletion, and delivery; this paper uses observability feasibility as the service boundary, residual-energy rotation to delay depletion of informative nodes, and link-aware scoring only to evaluate delivery efficiency under loss.

III. Methodology

Fig. 1 summarizes the overall workflow of the proposed observability-guided scheduling framework. The framework first establishes an observability-based service boundary from the sensing model, then constructs a greedy minimum-feasible active set, and finally applies residual-energy-aware rotation to preserve informative sensing directions over time. The lossy-delivery component is evaluated separately as a scope check rather than being included in the primary scheduling constraint.

A. Field, sensing, and service model

The monitored field is represented by the finite-dimensional linear system

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{w}_k, \qquad \mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}), \tag{1}$$

where $\mathbf{x}_k \in \mathbb{R}^d$. The finite-dimensional state is an estimator-oriented discretization that permits explicit matrix tests; it is not a claim that every spatial field is intrinsically finite-dimensional. The synthetic field uses a grid and

$$\mathbf{A} = \alpha_s \mathbf{I} + \beta_s \mathbf{W}, \tag{2}$$

where $\mathbf{W}$ is a row-normalized four-neighbor coupling matrix without wrap-around boundaries. This stable spatially coupled instance is a reproducible benchmark, not a unique physical law.

Node i observes

$$y_{i,k} = \mathbf{H}_i \mathbf{x}_k + v_{i,k}, \qquad v_{i,k} \sim \mathcal{N}(0, r). \tag{3}$$

Nodes are randomly placed in the unit square. Their rows use a normalized local Gaussian profile over grid cells within radius ρ_h:

$$H_{i\ell} = \frac{\exp(-\|\mathbf{z}_\ell - \mathbf{p}_i\|_2^2/(2\sigma_h^2))}{\sum_{\ell' \in \mathcal{R}_i} \exp(-\|\mathbf{z}_{\ell'} - \mathbf{p}_i\|_2^2/(2\sigma_h^2))}, \quad \ell \in \mathcal{R}_i, \tag{4}$$

with $H_{i\ell} = 0$ otherwise and nearest-cell fallback for empty neighborhoods. The profile represents localized spatial influence rather than a particular transducer response; unit-sum normalization makes Gramian comparisons reflect geometry rather than arbitrary gain. For active set $\mathcal{S}_k$, $\mathbf{C}_{\mathcal{S}_k}$ stacks the selected rows. Scheduling runs at a sink or edge controller using nominal estimates of $\mathbf{A}$ and $\{\mathbf{H}_i\}$, together with periodic residual-energy reports; model-free operation is outside scope.

B. Observability boundary and scheduling

For window T_o, reconstructability is measured by

$$\mathbf{W}_o(\mathcal{S}) = \sum_{t=0}^{T_o-1} (\mathbf{A}^t)^\top \mathbf{C}_{\mathcal{S}}^\top \mathbf{C}_{\mathcal{S}} \mathbf{A}^t, \tag{5}$$

$$\lambda_{\text{obs}}(\mathcal{S}) = \lambda_{\min}(\mathbf{W}_o(\mathcal{S})). \tag{6}$$

The full network sets a reference margin and threshold

$$\lambda_{\text{full}} = \lambda_{\text{obs}}(\mathcal{V}), \qquad \varepsilon = \eta\lambda_{\text{full}}. \tag{7}$$

Here η is a dimensionless, deployment-calibrated service margin rather than a reliability probability. The main setting uses $\eta = 0.15$ from the tested 0.10–0.25 range.

Algorithm 1 Observability-feasible residual-energy-aware scheduling

Require: $\mathbf{A}$, sensing rows $\{\mathbf{H}_i\}$, T_o, η, T_r, energies, and optional link probabilities $\{p_i\}$.
Ensure: Active sets $\{\mathcal{S}_k\}$, delivered sets $\{\mathcal{D}_k\}$ under loss, and metrics.
Compute $\varepsilon \leftarrow \eta\lambda_{\text{obs}}(\mathcal{V})$; set $\mathcal{S}_{-1} \leftarrow \emptyset$.
for $k = 0, \ldots, K-1$ **do**
  $\mathcal{A}_k \leftarrow \{i : E_i(k) > 0\}$.
  if $k = 0$ or $k \bmod T_r = 0$ or $\mathcal{S}_{k-1} \notin \mathcal{A}_k$ **then**
    Set $\mathcal{S}_k \leftarrow \emptyset$.
    while $\lambda_{\text{obs}}(\mathcal{S}_k) < \varepsilon$ and $\mathcal{A}_k \setminus \mathcal{S}_k \neq \emptyset$ **do**
      for all $j \in \mathcal{A}_k \setminus \mathcal{S}_k$ **do**
        Compute Δ_j and $s_j \leftarrow \alpha_r \,\text{norm}(\Delta_j) + (1-\alpha_r)\,\text{norm}(E_j(k))$.
        if link-aware mode **then**
          $s_j \leftarrow \alpha_o \,\text{norm}(\Delta_j) + \alpha_e \,\text{norm}(E_j(k)) + \alpha_c \,\text{norm}(p_j)$.
        end if
      end for
      Add the highest-scoring candidate to $\mathcal{S}_k$.
    end while
  else
    $\mathcal{S}_k \leftarrow \mathcal{S}_{k-1} \cap \mathcal{A}_k$.
  end if
  Collect from $\mathcal{S}_k$; under loss form $\mathcal{D}_k \subseteq \mathcal{S}_k$.
  Update active and sleep energy costs.
end for

For fixed-ratio pruning with keep ratio r_{keep}, define $m = \lceil r_{\text{keep}} N \rceil$, where $N = |\mathcal{V}|$. The greedy candidate is

$$j^\star = \arg\max_{j \in \mathcal{V}\setminus\mathcal{S}} \lambda_{\text{obs}}(\mathcal{S} \cup \{j\}). \tag{8}$$

Equation (8) is repeated until $|\mathcal{S}| = m$. For long-term scheduling, nodes are added by marginal gain

$$\Delta_j = \lambda_{\text{obs}}(\mathcal{S} \cup \{j\}) - \lambda_{\text{obs}}(\mathcal{S}) \tag{9}$$

until $\lambda_{\text{obs}}(\mathcal{S}) \geq \varepsilon$. This is a greedy minimum-feasible construction, not a global minimum-cardinality guarantee. Every T_r steps or after node death, surviving candidates are scored by

$$\text{score}(j) = \alpha_r \,\text{norm}(\Delta_j) + (1-\alpha_r)\,\text{norm}(E_j), \tag{10}$$

where E_j is residual energy. Here, $\text{norm}(\cdot)$ denotes min–max normalization over the current candidate set, with zero range mapped to zero. Energies evolve as

$$E_i(k+1) = \begin{cases} E_i(k) - e_{\text{act}}, & i \in \mathcal{S}_k, \\ E_i(k) - e_{\text{sleep}}, & i \notin \mathcal{S}_k. \end{cases} \tag{11}$$

The same normalization definition is used in Algorithm 1.

The motivation of residual-energy-aware rotation is that observability contribution is generally non-uniform across sensors. A node with a larger marginal observability gain Δ_j is more likely to be repeatedly selected by a purely observability-driven scheduler. Although such selection improves short-term reconstruction feasibility, repeatedly activating the same informative nodes accelerates their depletion and may create future sensing-direction bottlenecks. Therefore, the energy term does not replace observability-based selection; instead, it regulates the temporal usage of informative sensors. By redistributing activation load among nodes with comparable observability contributions, the scheduler preserves feasible sensing geometry for a longer service horizon.

C. Delivery extension and metrics

For delivery-efficiency analysis, $p_i = \exp(-\gamma d_i)$ is clipped to $[0.05, 0.99]$ and added to the rotation score, where d_i is the node-to-sink distance and γ controls link-decay severity. Filtering uses only delivered rows. State reconstruction uses a standard Kalman filter with the measurement rows available in $\mathcal{S}_k$ under ideal delivery and in $\mathcal{D}_k$ under lossy delivery. This extension is not chance-constrained and does not guarantee post-delivery observability; it diagnoses the efficiency loss caused by realized packet outcomes. Reconstruction accuracy is measured by

$$\text{RMSE} = \sqrt{\frac{1}{Kd}\sum_{k=1}^{K} \|\mathbf{x}_k - \hat{\mathbf{x}}_k\|_2^2}. \tag{12}$$

Active ratio, T_{obs} (the first time step at which $\lambda_{\text{obs}} < \varepsilon$), T_{FND}, energy dispersion, and delivery metrics are also reported. Selection and re-selection cost $\mathcal{O}(mNT_o d^3)$ in the tested centralized implementation.

IV. Experiments

A. Setup and evaluation protocol

We evaluate three ways in which reconstruction service can fail: loss of sensing geometry, depletion of informative nodes, and packet loss. Unless varied explicitly, the synthetic setting uses the model in Section III with $d = 16$, $T_o = 6$, $q = r = 0.005$, and $\eta = 0.15$. Long-term runs use $N = 80$, $K = 160$, $E_0 = 1$, $e_{\text{act}} = 0.015$, $e_{\text{sleep}} = 0.002$, $\alpha_r = 0.7$, and $T_r = 10$. The communication score uses $(\alpha_o, \alpha_e, \alpha_c) = (0.55, 0.25, 0.20)$ and $p_i = \exp(-\gamma d_i)$.

The default η and score weights are balanced operating points from the tested sweeps, not universal constants or single-metric optima. Baselines use the same nominal activation budget and rescheduling period where applicable. Results are means with 95% confidence intervals over 10 deployments for pruning and long-term comparisons and six deployments for deterministic sensitivity, communication, and trace checks. The Intel Berkeley Lab trace is binned into 10-minute temperatures, projected to six PCA dimensions, and fitted with a least-squares AR(1) model.

B. Structural pruning and scale

Fixed-ratio pruning compares observability-driven, random, coverage-based, and energy-aware selection. At $r_{keep} = 0.5$, observability-driven pruning remains feasible with $\lambda_{\text{obs}} \approx 0.00621$, versus 0.00373 for coverage-based pruning; random and energy-aware selection fall below the full-network margin. Fig. 2 summarizes the scaling result: at $\eta = 0.15$, the minimum active ratio falls from 0.467 at $N = 30$ to 0.179 at $N = 100$, while the threshold sweep exposes the active-ratio/service-margin trade-off. Across the tested $\eta = 0.10$–0.25 range, the active ratio rises from 0.226 to 0.345 while T_{obs} remains between 130.7 and 160 steps. Thus additional nodes help only when they add useful sensing directions.

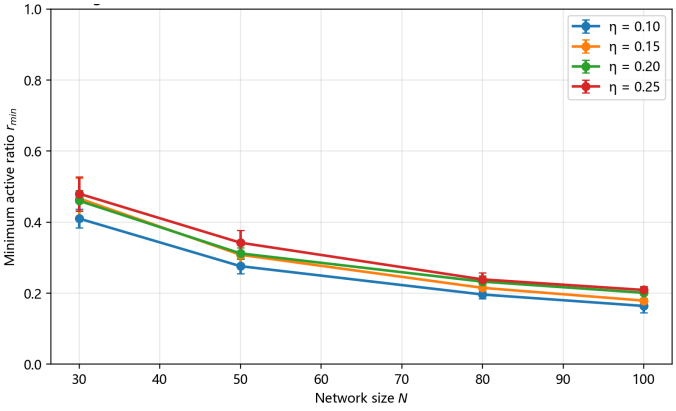


Fig. 2: Minimum active ratio required to satisfy the observability threshold as network size and threshold ratio vary.

C. Long-term service and mechanism

Table I compares the long-term baselines at $N = 80$. Obs+Rotation reduces the horizon-averaged active ratio to 0.264, gives RMSE 0.1005, and extends T_{obs} to 158.5 steps. Full and Obs-Only both reach $T_{\text{obs}} = 67.0$; energy-only and round-robin methods can reach $T_{\text{FND}} \approx 160$ but violate the structural threshold initially. Coverage+Energy and KF-Greedy Rotation have competitive RMSE but short service lifetimes, while KF-Greedy+Repair confirms that explicit observability repair restores the lifetime boundary. Obs+Rotation satisfies the scheduled threshold in 99.1% of rounds with mean normalized margin 1.127.

Removing rotation lowers T_{obs} from 158.5 to 67.0 and raises RMSE to 0.1440; removing the energy term shortens lifetime to 142.5; removing observability gives $T_{\text{obs}} = 0$. The same pattern holds for a static observable subset. Noise and failure checks remain consistent: at $r = 0.02$, Obs+Rotation has RMSE 0.1183 versus 0.1277 for Full, and with 40% initial failures it retains $T_{\text{obs}} = 119.2$. The tested implementation takes 0.115 ± 0.010 s at $N = 80$, with centralized selection cost $\mathcal{O}(mNT_o d^3)$.

D. Lossy delivery and scope

Table II separates scheduled activation from delivered observations. Relative to Obs+Rotation with loss, Comm-Obs Rotation increases useful observations from 12.2 to 13.5 per round and reduces energy per useful observation from 0.034 to 0.032, while RMSE remains 0.112 and strict delivered T_{obs} remains zero. Adding 20% redundancy lowers RMSE to 0.101 and increases useful observations to 17.9, at the cost of more transmissions. As γ increases from 0.5 to 2.0, PDR falls from 0.828 to 0.489 and RMSE rises from 0.1030 to 0.1131. These results characterize delivery efficiency, not a chance-constrained post-delivery guarantee.

The Intel trace check gives RMSE 0.442, $T_{\text{obs}} = 120$, and active ratio 0.286 for Obs+Rotation, versus RMSE 1.224 for Full. It uses measured connectivity but not a complete MAC trace; deployment-level protocol validation remains outside scope.

V. Results and Discussion

The results support three linked findings. First, geometric redundancy is not equivalent to reconstruction redundancy. Observability-driven pruning retains the structural margin at aggressive retention, and the scaling curve in Fig. 2 shows that the required active ratio decreases with network density but remains deployment-dependent.

Second, long-term service is different from hardware lifetime. Table I shows that energy-only and round-robin baselines can delay first node death while violating the observability boundary from the first round. Obs+Rotation instead keeps the scheduled set feasible for 158.5 steps with RMSE 0.1005 and mean active ratio 0.264. Its advantage comes from role rotation: removing rotation reduces T_{obs} to 67.0, while removing observability removes service entirely. The stronger KF-Greedy+Repair diagnostic reaches a similar lifetime, indicating that the result is due to explicit feasibility preservation rather than the particular Gramian scoring primitive.

The mechanism behind this lifetime extension is the uneven contribution of individual nodes to the observability Gramian. A node with a larger marginal observability gain

$$\Delta_j = \lambda_{\text{obs}}(S \cup \{j\}) - \lambda_{\text{obs}}(S)$$

is more likely to be repeatedly selected by a purely observability-driven scheduler. Although this improves short-term reconstruction feasibility, repeated activation can exhaust structurally important nodes and create sensing-direction bottlenecks. Residual-energy rotation therefore does not simply reduce energy consumption; it redistributes the depletion of informative sensors over time and preserves feasible sensing geometry for a longer horizon. In this sense, the benefit of rotation comes from protecting information contribution balance rather than only extending battery lifetime.

RMSE and service lifetime must therefore be read together. A stable filter can retain a moderate average error after structural observability has been lost, whereas the threshold metric records the loss of guaranteed sensing geometry. The 99.1% scheduled-observability rate and normalized margin 1.127 for Obs+Rotation make this distinction explicit.

Third, packet delivery is a separate boundary. Table II shows that link-aware scoring improves useful observations and energy efficiency, while limited redundancy improves RMSE at higher transmission cost. Strict delivered T_{obs} remains zero for lossy methods, confirming that the extension is a delivery-efficiency analysis rather than a chance-constrained scheduler. The trace check and alternative link models preserve the qualitative ordering, but do not replace protocol-stack validation.

VI. Limitations and Scope

The study uses a finite-dimensional linear approximation with approximate model knowledge; nonlinear or strongly time-varying fields may require a different estimator, and model mismatch can change selected sensing directions. The threshold ratio η is an empirical, deployment-calibrated service margin. The centralized greedy cost $\mathcal{O}(mNT_o d^3)$ limits

TABLE I: Long-term scheduling results at $N = 80$, $K = 160$; means $\pm$ 95% CI.

Method	Active ratio	RMSE	T_{obs}	T_{FND}	σ_E
Full (all alive)	0.419 ± 0.000	0.1325 ± 0.0074	67.0 ± 0.0	67.0 ± 0.0	0.0000 ± 0.0000
Obs-Only	0.092 ± 0.005	0.1440 ± 0.0049	67.0 ± 0.0	67.0 ± 0.0	0.3595 ± 0.0069
Obs+Rotation	0.264 ± 0.010	0.1005 ± 0.0024	158.5 ± 1.5	92.0 ± 8.2	0.1244 ± 0.0071
Energy-only	0.219 ± 0.012	0.1024 ± 0.0019	$0.0^{\dagger} \pm 0.0$	160.0 ± 0.0	0.0463 ± 0.0147
Duty-Cycle RR	0.203 ± 0.011	0.1069 ± 0.0020	$0.0^{\dagger} \pm 0.0$	160.0 ± 0.0	0.0486 ± 0.0101
Residual-Energy RR	0.205 ± 0.011	0.1081 ± 0.0019	$0.0^{\dagger} \pm 0.0$	150.4 ± 4.2	0.1639 ± 0.0056
Coverage+Energy	0.219 ± 0.012	0.0993 ± 0.0011	1.0 ± 1.9	160.0 ± 0.0	0.0535 ± 0.0073
KF-Greedy Rotation	0.219 ± 0.012	0.0971 ± 0.0020	8.0 ± 7.2	122.0 ± 17.8	0.1342 ± 0.0146
KF-Greedy+Repair	0.252 ± 0.012	0.0970 ± 0.0022	158.8 ± 1.2	94.7 ± 9.8	0.1474 ± 0.0092

†Initial strict-threshold violation.

TABLE II: Lossy-delivery results; means $\pm$ 95% CI.

Method	RMSE	PDR	Active tx	Useful obs	Energy/useful
Obs+Rotation (Ideal)	0.101 ± 0.004	1.00 ± 0.00	19.6 ± 1.1	19.6 ± 1.1	0.021 ± 0.000
PDR+Energy Rotation	0.118 ± 0.005	0.66 ± 0.01	17.2 ± 0.9	11.4 ± 0.5	0.034 ± 0.001
Obs+Rotation + Loss	0.112 ± 0.004	0.62 ± 0.02	19.6 ± 1.1	12.2 ± 0.6	0.034 ± 0.001
Comm-Obs Rotation	0.112 ± 0.002	0.64 ± 0.01	21.1 ± 0.9	13.5 ± 0.6	0.032 ± 0.001
Comm-Obs + 20% Redund.	0.101 ± 0.001	0.64 ± 0.02	28.1 ± 2.3	17.9 ± 1.5	0.029 ± 0.001

direct scaling without cached updates, candidate screening, or distributed coordination. Packet loss is modeled by node-level success probabilities rather than a full MAC trace, so no strict post-delivery guarantee or deployment-level reliability claim is made.

VII. Conclusions

This paper studied node scheduling for energy-constrained WSNs from a reconstruction-service perspective. The contribution is not a new sensor-selection primitive, but a scheduling formulation that treats observability preservation as the lifetime objective. Keeping sensors alive is not equivalent to keeping the field reconstructable. Observability feasibility provides the service boundary used in this study, and observability lifetime makes this loss of service explicit.

The proposed scheduling framework combines observability-constrained activation with residual-energy-aware rotation. In the representative 80-node long-term setting, it reduces active ratio and RMSE while extending T_{obs} from 67.0 to 158.5 steps. The rotation mechanism addresses the early depletion of structurally informative sensors. Energy-only and round-robin baselines can delay node death, but their active sets still remain below the strict structural observability margin.

The lossy-delivery experiments show that scheduled observability does not automatically translate into delivered observability. Link-aware scoring and redundancy improve delivered useful observations. Strict post-delivery guarantees still require expected-observability, chance-constrained scheduling, or retransmission-aware redundancy. These directions remain future work. The present evidence is based on controlled simulations and trace-driven checks; deployment-level validation remains outside the scope of this paper.

References

[1] D. Lin, Q. Wang, W. Min, J. Xu, and Z. Zhang, "A survey on energy-efficient strategies in static wireless sensor networks," *ACM Trans. Sen. Netw.*, vol. 17, no. 1, 2020, art. no. 3.

[2] H. Zhang, R. Ayoub, and S. Sundaram, "Sensor selection for kalman filtering of linear dynamical systems: Complexity, limitations and greedy algorithms," *Automatica*, vol. 78, pp. 202–210, 2017.

[3] K. Yamada, Y. Sasaki, T. Nagata, K. Nakai, D. Tsubakino, and T. Nonomura, "Efficient sensor node selection for observability gramian optimization," *Sensors*, vol. 23, no. 13, 2023, art. no. 5961.

[4] S. D. Bopardikar, "A randomized approach to sensor placement with observability assurance," *Automatica*, vol. 123, 2021, art. no. 109340.

[5] H. Yetgin, K. T. K. Cheung, M. El-Hajjar, and L. H. Hanzo, "A survey of network lifetime maximization techniques in wireless sensor networks," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 828–854, 2017.

[6] Y. Mo, R. Ambrosino, and B. Sinopoli, "Sensor selection strategies for state estimation in energy constrained wireless sensor networks," *Automatica*, vol. 47, no. 7, pp. 1330–1338, 2011.

[7] A. Hashemi, M. Ghasemi, H. Vikalo, and U. Topcu, "Randomized greedy sensor selection: Leveraging weak submodularity," *IEEE Transactions on Automatic Control*, vol. 66, no. 1, pp. 199–212, 2021.

[8] R. Vafaee and M. Siami, "Real-time sensor selection for time-varying networks with guaranteed performance," *Automatica*, vol. 163, 2024, art. no. 111550.

[9] J. Abou Rahal and G. de Veciana, "Scheduling "last minute" updates for timely decision-making," in *Proceedings of the Int'l ACM Conference on Modeling Analysis and Simulation of Wireless and Mobile Systems*, 2023, pp. 65–72.