Journal of Emerging Paradigms in Computing Systems (JEPCS) | Volume 04, Issue No. 01, 2026

ARTICLE TYPE: Research Article

Article ID: JEPCS-2026-0401-0003

Machine Learning-Assisted Reliability-Aware Digital Twin Synchronization for Resource-Constrained IoT Sensor Networks

Zupash Awais [1,*], **Nauman Irshad Ali Shah** [1]

[1] Faculty of Information Technology & Computer Science, University of Central Punjab

* Corresponding author: zupash.awais@ucp.edu.pk

Received: April 02, 2026 **Accepted:** June 22, 2026

Abstract- Continuous synchronization allows a Digital Twin (DT) to maintain a useful representation of a physical Internet of Things (IoT) device, but frequent updates consume wireless resources and do not guarantee freshness when some links are unreliable. This paper presents a reliability-aware DT synchronization method for resource-constrained sensor networks. A lightweight autoregressive DT estimates temperature between successful updates. The update scheduler then combines prediction behavior, Age of Information (AoI), measured link reliability, and radio cost. For sensors with weak direct connectivity, a two-hop path is selected only when its expected delivery probability per radio transmission is better than direct delivery. The evaluation uses a five-day Intel Berkeley Research Lab temperature trace with 52 nodes and the corresponding aggregate asymmetric connectivity measurements. A chronological 60/20/20 train-validation-test protocol is used, and the main held-out comparison is repeated over 30 matched wireless-delivery realizations. At five radio transmissions per 30-second slot, equal to 9.62% of full telemetry, the proposed method obtains 0.2076 ± 0.0042 RMSE and 12.00-slot mean AoI. The direct-link-aware DT obtains 0.3921 ± 0.0330 and 32.01 slots under the same budget. Thus, the proposed method reduces RMSE by 47.07% and mean AoI by 62.52% relative to this strongest direct-link ablation. The results support reliability-aware selective synchronization when radio capacity is limited.

Keywords- Digital Twin; Internet of Things; wireless sensor networks; Age of Information; reliability-aware synchronization

1. Introduction

Internet of Things (IoT) systems obtain physical-state information through distributed sensing devices and deliver these measurements to monitoring, control, and decision-support applications. In many deployments, sensing is inexpensive compared with radio communication. The sensor node may also have limited energy, memory, and processing capacity. When every measurement is forwarded, communication resources are spent even when the physical state changes slowly. The opposite choice, aggressive suppression, can leave the receiver with an outdated view of the monitored process.

DT technology changes this trade-off because the receiver can maintain a virtual state between physical updates. Recent surveys describe DTs as an increasingly important component of wireless and networked systems, with synchronization quality, model fidelity, and communication overhead treated as coupled design issues [1]-[4]. A DT is therefore not useful only because it contains a model. Its state must be corrected often enough to remain consistent with the physical entity, and those corrections must be delivered over the available network.

The synchronization problem becomes more difficult when link quality differs across devices. A sensor with a stable physical signal may require few updates, while another sensor may need frequent corrections but have a weak direct path to the collection point. Recent work has addressed DT synchronization through continual reinforcement learning, hybrid inverse reinforcement learning, semantic communication, wireless resource allocation, and explicit synchronization budgets [5]-[9]. These methods show that selecting when and what to synchronize is a resource-management problem. Most, however, evaluate abstract communication resources or optimized scheduling models rather than the measured asymmetric link structure of a real environmental sensor deployment.

Freshness adds a second dimension to reconstruction accuracy. The 2025 review by Loubany et al. [10] summarizes synchronization metrics and introduces AoI as a DT-oriented extension of freshness measures. Recent papers also use AoI, statistical AoI, state staleness, or related fidelity measures when deciding when a DT should migrate, refresh, or synchronize [11]-[15]. These metrics are relevant because a numerically plausible estimate may still be based on an old physical update. In a resource-constrained network, the scheduler must therefore consider both the error expected from skipping an update and the age of the last successful observation.

The present paper considers this problem using the Intel Berkeley Research Lab sensor deployment. Recent papers continue to reuse this trace for communication-efficient sensing, as discussed in Section 2.4. Unlike papers that use only the temperature series, the present evaluation also uses the released aggregate connectivity probabilities. This permits a trace-driven question that is usually abstracted away: if a node has a poor direct link, should the system spend scarce radio resources on repeated direct attempts, retain the DT estimate, or select a two-hop path through a better-connected node?

To answer this question, a lightweight DT is combined with a reliability-aware update scheduler. Each sensor has an autoregressive virtual-state predictor. At every 30-second slot, the scheduler ranks candidate updates from their prediction behavior, current AoI, route reliability, and radio cost. Direct and two-hop paths are compared using the recorded connectivity values, and relaying is used only when the expected delivery probability per radio transmission improves. The approach is intentionally compact so that the contribution of each component can be isolated through ablation rather than hidden inside a large learned policy.

The major contributions of this paper are summarized below:

- The synchronization policy assigns each sensor an update score based on its prediction behavior, AoI, recorded path reliability, and transmission cost. Updates are then selected subject to a fixed radio budget in each time slot.
- For every scheduled sensor, the direct route is compared with its best available two-hop alternative using the asymmetric Intel connectivity matrix. A relay is selected only when the two-hop delivery probability, after accounting for the cost of two radio transmissions, is better than direct delivery.
- The method is evaluated on a real five-day temperature trace containing 52 sensor nodes. Training, validation, and testing are performed in chronological order, and the held-out evaluation is repeated over 30 matched wireless realizations.
- At 9.62% of full radio traffic, the proposed method obtains 0.2076 ± 0.0042 RMSE, compared with 0.3921 ± 0.0330 for the strongest direct-link DT ablation. Mean AoI is reduced from 32.01 to 12.00 slots. The proposed method also retains its accuracy advantage across the tested 3.85%-28.85% traffic range.

The remainder of the paper is organized as follows. Section 2 reviews recent work on DT synchronization, freshness-aware state maintenance, edge-resource management, and adaptive sensing. Section 3 presents the system model and the proposed synchronization method. Section 4 describes the dataset, experimental protocol, baselines, and modeling assumptions. Section 5 presents the results and limitations. Section 6 concludes the paper and highlight future work.

2. Related Work

The synchronization problem considered in this paper is related to four areas of recent research: DT synchronization and wireless resource scheduling, freshness-aware DT maintenance, edge-resource management, and communication-efficient sensing. The first two determine when a physical state should update its virtual counterpart, whereas the latter two address the cost of maintaining that state under limited network resources. Table 1 summarizes the papers most closely related to these aspects.

2.1. Digital Twin Synchronization and Wireless Resource Scheduling

Several recent papers formulate DT synchronization as a resource allocation problem rather than a fixed periodic update process. Tong et al. [5] select physical devices and wireless resource blocks through continual reinforcement learning so that the synchronization policy can adapt to changes in available capacity. Zhang et al. [6] consider an industrial setting and use hybrid inverse reinforcement learning to derive scheduling decisions from historical behavior. Synchronization delay is addressed by Li et al. [7] through joint semantic communication and mobile-edge resource allocation. Yu et al. [8] jointly optimize wireless resources and network synchronization, while Li et al. [9] impose an explicit synchronization budget for fidelity-aware edge queries.

A common feature of these approaches is that synchronization decisions are made under communication or computational constraints. The network condition is usually represented through quantities such as allocated resource blocks, transmission power, delay, or aggregate capacity. The present work considers a different aspect of the same problem. Its radio budget is combined with the measured asymmetric probability that one sensor node can successfully deliver an update to another node. Route reliability therefore directly affects whether a physical measurement is selected for DT correction.

Broader Network DT research provides the architectural context for this type of closed-loop synchronization. Bhatia and Kumar [1] review DT technology for wireless communication, while Tran et al. [2] describe Network DT architectures for multi-domain 6G ecosystems. Li et al. [3] examine the use of generative AI in Network DTs, and Sanjalawe et al. [4] review AI-driven DT networks for future wireless systems. These papers emphasize state exchange, predictive modeling, and closed-loop network control. In the resource-constrained sensor setting considered here, the corresponding decision is narrower: the system must determine which physical measurements should correct the DT when only a small number of radio transmissions are available and direct links have different success probabilities.

2.2. Freshness, Staleness, and Digital Twin Fidelity

Freshness is increasingly treated as part of DT synchronization rather than as a separate networking metric. Loubany et al. [10] review synchronization metrics for IoT networks and distinguish conventional AoI from DT-oriented measures that also account for physical-virtual consistency. Noroozi et al. [11] include AoI in DT migration decisions, linking migration cost with the age of the available state. Xiao et al. [12] introduce Statistical AoI as a risk-aware status-update metric that captures behavior not represented by average age alone. At the network level, Wang et al. [13] use a dual-time-scale slicing strategy for IIoT DT synchronization, whereas Li et al. [14] directly optimize DT freshness in edge computing.

Freshness is also tied to the quality of services that depend on the virtual state. Khalaf et al. [15] use UAV assistance to maintain DT accuracy and synchronization in IoT environments. Zhang et al. [16] paper mobility-aware DT migration in edge computing. Ai et al. [17] consider fidelity-aware inference together with service-model retraining, while Li et al. [18] optimize inference-service fidelity in DT-assisted edge computing. Ai et al. [19] jointly optimize model retraining and inference, making data drift and model refresh part of the DT service process. Shi et al. [20] address real-time spatiotemporal synchronization in a human-robot DT through adaptive data processing. Across these formulations, synchronization quality is treated as a graded property that depends on when the physical state is refreshed and how accurately the virtual system represents it.

The present evaluation therefore reports reconstruction error and state age together. RMSE and MAE quantify the numerical mismatch between the physical and virtual temperature states, while mean and 95th-percentile AoI measure how long the collection point has operated without a successful physical correction. These metrics make it possible to distinguish a numerically accurate estimate from a virtual state that has remained uncorrected for a long period. They are not intended to replace richer Age of Digital Twin (AoDT) or service-fidelity measures, but they provide a direct interpretation for the sensor trace used in this paper.

2.3. Resource Efficiency, Edge Intelligence, and Distributed DT Operation

Maintaining a DT also consumes communication, computation, and energy resources that are shared with other network functions. Ndikumana et al. [21] use DT-supported closed loops for energy-aware Open RAN fixed wireless access. Qadir et al. [22] paper DT-assisted multi-layer networks with low-latency and energy-efficient communication objectives. Chu et al. [23] develop asynchronous federated learning for a DT-empowered IoT network with explicit resource scheduling, whereas Zhou et al. [24] jointly consider federated learning, migration, and energy efficiency in DT edge networks.

The same resource limitation appears in application-oriented DT systems. Kwon [25] combines DT operation with IoT energy management for smart manufacturing. Chang et al. [26] address synchronized data transmission and dynamic device management in DT-enabled IIoT through a blockchain-based protocol. Yahya et al. [27] optimize task offloading for UAV-IoT DT networks under spatial uncertainty, while Hayawi et al. [28] paper DT-assisted task offloading for fog-node workload management.

These papers mainly allocate computation, migration, workload, or energy resources. The present problem concerns the physical-state update itself. The limited resource is the number of radio transmissions available in each time slot. For each sensor, the scheduler must decide whether the state should be transmitted directly, sent through a relay, or temporarily maintained using the DT prediction. This formulation allows the effects of temporal prediction, state freshness, measured link reliability, and transmission cost to be evaluated within the same synchronization process. Although the application settings differ, the common requirement is that virtual-state quality depends on when new physical information arrives and on the cost of obtaining it.

The proposed method is smaller than these end-to-end frameworks. It does not optimize computation offloading, federated model training, or blockchain synchronization. Its purpose is to isolate the communication layer of virtual-state maintenance in a sensor network. The resulting scheduler can therefore be evaluated with a small number of components and a direct ablation of temporal prediction, link awareness, and two-hop recovery.

2.4. Communication-Efficient Sensing

Communication reduction remains an active WSN problem even as DT methods become more common. Alanazi and Alanazi [29] use the Intel Lab data in a 2025 aggregation-aware MAC paper and predict transmission probability for temperature, humidity, light, and voltage. Lu et al. [30] use the Intel Berkeley temperature data in a 2026 adaptive-sampling framework where Kalman-filter uncertainty controls sparse sampling and sentinel nodes increase sampling around detected anomalies. Kangethe et al. [31] propose a volatility-aware predictive communication framework in 2026 and report that local prediction residuals can be used to avoid redundant IoT transmissions.

These papers are close to the data-reduction side of the present problem, but their objectives differ. The aggregation-aware work does not maintain an explicit DT freshness state. The sentinel framework emphasizes anomaly visibility, and the volatility-aware method focuses on residual-triggered communication. The present paper instead keeps a virtual state for every node, tracks the age of the last successful physical correction, and uses the released Intel connectivity probabilities to decide whether the update should be sent directly or through a selected relay. This difference is important because a sensor selected by a predictive trigger can still fail repeatedly when its direct path is poor.

Table 1. *Comparison with recent work.*

Paper	Year	Primary objective	Decision signal / resource	Difference from this paper
Tong et al. [5]	2025	Continual-RL DT synchronization	Device and resource-block selection	Abstract resource capacity; no measured WSN links
Zhang et al. [6]	2025	IIoT DT scheduling	Hybrid inverse RL	No trace-driven weak-link recovery
Li et al. [7]	2025	Delay-aware DT synchronization	Semantic communication + MEC resources	Latency objective; no measured node topology
Yu et al. [8]	2025	Network DT synchronization	Wireless resource management	Network resource model rather than Intel link matrix
Li et al. [9]	2025	Budget-constrained DT updates	Explicit update budget and staleness	Query fidelity focus; no route reliability
Loubany et al. [10]	2025	Synchronization metrics	AoI/AoDT taxonomy	Metric review rather than route-selection mechanism
Noroozi et al. [11]	2025	DT migration	AoI-aware migration	Mobility/migration rather than sensor reconstruction
Wang et al. [13]	2026	IIoT synchronization	Dual-time-scale slicing	No released WSN connectivity trace
Li et al. [14]	2026	DT freshness	Freshness/cost optimization	Edge-service model; no weak-link relay ablation
Khalaf et al. [15]	2026	UAV-aided IoT DT	Accuracy and synchronization	UAV-assisted framework rather than fixed sensor topology
Qadir et al. [22]	2025	DT multi-layer networking	Latency and energy efficiency	System-layer optimization, not node-level AoI
Chu et al. [23]	2026	DT-empowered IoT learning	Asynchronous FL resource scheduling	Learning-resource focus rather than state reconstruction
Zhou et al. [24]	2026	DT edge learning/migration	Energy-efficient FL and migration	No sensor-link reliability ablation
Zhang et al. [16]	2025	DT migration	Mobility-aware DRL	Migration decision rather than update routing
Ai et al. [17]	2025	Inference fidelity	Service-model retraining	Compute/model fidelity focus
Chang et al. [26]	2025	IIoT synchronized transmission	Clock sync + device management	Security/device-management focus
Alanazi and Alanazi [29]	2025	WSN data transmission	Intel Lab + adaptive aggregation	No DT AoI or route-cost decision
Lu et al. [30]	2026	Adaptive anomaly sampling	Intel temperature + sentinel wake-up	Anomaly-window objective; no DT route selection
Kangethe et al. [31]	2026	Predictive IoT transmission	Volatility-aware residual trigger	No network-reliability-aware DT state
This paper	2026	Reliability-aware DT synchronization	Real temperature trace + connectivity + AoI + cost	Temperature only; aggregate link probabilities

3. Proposed System

3.1. System Model and Design Objective

The considered network contains $N = 52$ active sensor nodes and one collection endpoint. node i observes a physical temperature state $x_i(t)$ at discrete 30-second slot t. The collection point maintains the corresponding DT estimate $\hat{x}_i(t)$. A successful physical update replaces the current virtual estimate with the measured temperature. If the sensor is not scheduled or the packet is lost, the DT advances from its own prior state.

A radio budget B limits the number of individual transmissions that can be used in one slot. A direct update costs one transmission. A two-hop update through relay j costs two transmissions. Full direct telemetry therefore costs 52 transmissions per slot, whereas the main experiment uses $B = 5$. The objective is to keep the set of virtual states accurate and recent without exceeding this radio budget. This objective requires a joint view of physical-state predictability, time since the last successful update, route reliability, and route cost.

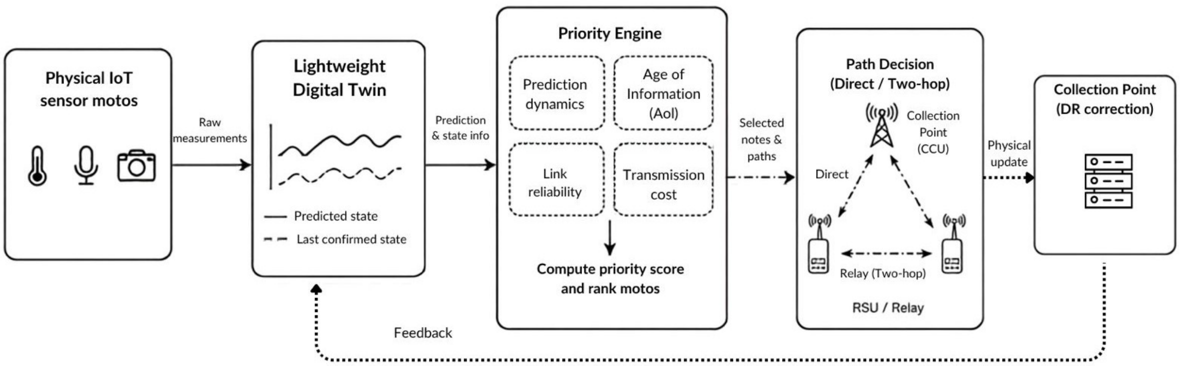

Figure 1: Reliability-aware DT synchronization workflow.

3.2. Lightweight Temporal Digital Twin

A separate second-order autoregressive model is fitted for each sensor using only the chronological training period. Ridge regularization with $\alpha = 0.1$ is used to reduce coefficient instability. Let $\beta_{i,0}$, $\beta_{i,1}$, and $\beta_{i,2}$ denote the fitted coefficients for sensor i. Its one-step autoregressive prediction is

$$\widehat{x}_i^{\mathrm{AR}}(t) = \beta_{i,0} + \beta_{i,1}\widehat{x}_i(t-1) + \beta_{i,2}\widehat{x}_i(t-2) \tag{1}$$

The virtual state is not forced to the full autoregressive prediction after every missed update. A damping factor γ controls how strongly the prior state moves toward the AR prediction:

$$\widetilde{x}_i(t) = \widehat{x}_i(t-1) + \gamma\left[\widehat{x}_i^{\mathrm{AR}}(t) - \widehat{x}_i(t-1)\right] \tag{2}$$

The training residual scale σ_i is the root mean squared residual of the fitted model. For scheduling, this value is normalized by the median residual scale across sensors and clipped to $[0.5, 3.0]$. The term gives more weight to sensors whose state is intrinsically harder to predict and avoids assigning identical priority to nodes that have very different prediction behavior.

3.3. Link Reliability and Cost-Aware Path Selection

The Intel connectivity file records an aggregate asymmetric success probability p_{ij} for a packet sent by node i and heard by receiver j. Receiver identifier 0 is treated as the collection endpoint in this evaluation. Thus, the direct success probability for sensor i is $p_{i,0}$. For each possible relay j, the modeled two-hop delivery probability is

$$r_{i,j}^{(2)} = p_{i,j}p_{j,0} \tag{3}$$

The two-hop outcomes are sampled independently in the trace-driven Monte Carlo evaluation. This is a modeling assumption because the released file contains aggregate link probabilities rather than time-stamped packet outcomes. The best relay is the active node with the largest $r_{i,j}^{(2)}$. A relay is used only when its expected success probability per radio transmission exceeds the direct alternative:

$$\text{use relay } j \quad \text{if} \quad \frac{r_{i,j}^{(2)}}{2} > p_{i,0} \tag{4}$$

The division by two charges the relay path for both radio transmissions. The rule is therefore not a shortest-path rule and does not select a relay merely because the two-hop success probability is larger. A path must provide enough reliability improvement to justify its additional radio cost.

3.4. Age- and Reliability-Aware Update Priority

The AoI $a_i(t)$ of a node is the number of 30-second slots since the collection point received its latest successful physical update. AoI increases after a suppressed or failed update and resets to zero after successful delivery. The DT first computes a normalized prediction-dynamics term

$$d_i(t) = \frac{|\widetilde{x}_i(t) - \widehat{x}_i(t-1)|}{\sigma_i + \varepsilon} \tag{5}$$

A base information value then combines state age, sensor-specific residual scale, and short-term prediction dynamics:

$$u_i(t) = [1 + w_A a_i(t)]\, \bar{\sigma}_i + 0.2 d_i(t) \tag{6}$$

For the proposed method, let r_i be the success probability of the selected direct or relay path and c_i be its radio cost, where c_i is 1 for direct transmission and 2 for a relay path. The final scheduling score is

$$s_i(t) = u_i(t)\frac{0.35 + 0.65 r_i}{c_i}\left[1 + 0.015 a_i(t)^{1.35}\right] \tag{7}$$

Candidate sensors are ordered by $s_i(t)$. Updates are admitted until the remaining radio budget cannot support the next path. The reliability-unaware temporal baseline removes the path-reliability factor and does not use relays. The direct-link-aware ablation uses $p_{i,0}$ in the priority score but restricts all updates to direct delivery. This construction separates the contribution of temporal prediction, link awareness, and cost-aware relaying. Algorithm 1 summarizes the per-slot operation.

Algorithm 1 Reliability-Aware DT Synchronization at Time Slot t

Require: Previous DT states, AoI vector, radio budget B, direct and relay probabilities
Ensure: Updated DT states, AoI vector, and radio usage

for each sensor i **do**
  Compute $\tilde{x}_i(t)$ using the damped AR(2) predictor.
  Update $a_i(t)$ and compute $d_i(t)$ and $u_i(t)$.
  Select the direct or precomputed two-hop candidate using the condition in Eq. (4).
  Compute $s_i(t)$ from Eq. (7).
end for
Sort sensors in descending order of priority $s_i(t)$.
for each sensor i in descending priority order **do**
  if the selected path cost fits the remaining radio budget **then**
    Admit the selected path.
    Sample the direct or two-hop transmission outcome.
    if delivery succeeds **then**
      Set $\hat{x}_i(t) \leftarrow x_i(t)$.
      Set $a_i(t) \leftarrow 0$.
    else
      Set $\hat{x}_i(t) \leftarrow \tilde{x}_i(t)$.
    end if
  end if
end for
return Updated DT states, AoI vector, and radio usage.

With N sensors, prediction and score construction require $O(N)$ work per slot. Sorting the candidate scores requires $O(N \log N)$, while packet-outcome sampling is bounded by the radio budget. The full connectivity matrix requires $O(N^2)$ storage, but route selection itself can be precomputed once for a static deployment. The online DT state, AoI, residual scale, and selected-route vectors require $O(N)$ additional memory. For the 52-node trace, these operations remain small compared with the sensing interval; the paper nevertheless reports them as algorithmic complexity rather than embedded-hardware latency.

3.5. Evaluation Metrics

Reconstruction quality is measured by root mean squared error (RMSE), mean absolute error (MAE), and the 95th percentile absolute error. For M evaluated physical-virtual pairs, RMSE and MAE are calculated as

$$\text{RMSE} = \sqrt{\frac{1}{M}\sum_{m}(x_m - \widehat{x}_m)^2}, \qquad \text{MAE} = \frac{1}{M}\sum_{m}|x_m - \widehat{x}_m| \tag{8}$$

Radio usage is expressed relative to full direct telemetry. If N_{TX} is the number of radio transmissions used by a method and N_{full} is the number required to attempt all sensor updates, communication reduction is

$$\text{Communication reduction} = 1 - \frac{N_{\text{TX}}}{N_{\text{full}}} \tag{9}$$

Delivery rate is the fraction of scheduled packet attempts that reach the endpoint. Freshness is reported with mean AoI and 95th-percentile AoI. The metrics are interpreted jointly. A low RMSE obtained by using almost full telemetry does not satisfy the communication objective, while a low traffic rate with very old physical corrections does not provide a useful synchronization policy.

Table 2. *Experimental parameters.*

Parameter	Value
Active nodes	52
Sampling interval	30 s
Training / validation / test	8,640 / 2,880 / 2,880
DT predictor	AR(2) Ridge
Main radio budget B	5 TX/slot
Budget sweep	2, 3, 5, 8, 10, 15 TX/slot
Main test repetitions	30
Budget sweep repetitions	10
Temporal DT (γ, w_A)	$(0.10, 0.40)$
Direct-link-aware DT (γ, w_A)	$(0.25, 0.80)$
Proposed (γ, w_A)	$(0.25, 0.80)$

4. Experimental Setup

4.1. Dataset and Preprocessing

The evaluation uses a preprocessed temperature trace derived from the Intel Berkeley Research Lab deployment [32]. The matrix contains 14,400 observations for 52 nodes over five consecutive days at 30-second resolution. nodes 5 and 15 are absent from this benchmark. The official release also contains node coordinates and aggregate asymmetric connectivity measurements, which are used to construct the trace-driven communication model rather than a synthetic topology. The implementation was executed in Python 3.13.5 with NumPy 2.3.5, pandas 2.2.3, scikit-learn 1.8.0, SciPy 1.17.0, and Matplotlib 3.10.8. Table 2 shows the experiment parameters.

Five zero-valued entries in the supplied temperature matrix are treated as missing observations and linearly interpolated within the corresponding sensor series. No future test observation is used for model fitting or parameter selection. The coordinates are retained for provenance, although the current path rule depends on measured delivery probability rather than Euclidean distance.

Table 3. *Distribution of direct-to-endpoint success probabilities across the 52 evaluated nodes.*

Direct success-probability range	nodes	Mean $p_{i,0}$	Relays selected
< 0.10	5	0.028	5
0.10-0.30	2	0.254	0
0.30-0.60	22	0.483	0
0.60-0.90	23	0.729	0
≥ 0.90	0	-	0

Table 3 shows why a reliability-aware comparison is needed. Most nodes are not equally difficult to reach, and the path rule selects relaying only for a small subset. The later node-level analysis therefore tests whether these weak links account for the difference between the direct-link-aware and proposed DT methods.

Table 4. *Main held-out results. Restricted methods use five radio transmissions per 30-second slot.*

Method	RMSE	MAE	P95 abs. err.	Comm. reduction	Delivery	Mean AoI
Full direct (lossy)	0.2764 ± 0.0661	0.0478	0.1407	0.00%	53.90%	10.49
Periodic	0.8176 ± 0.0677	0.2500	0.9517	90.38%	53.94%	69.81
Temporal DT	0.4069 ± 0.0392	0.2110	0.7984	90.38%	20.21%	35.46
DT + direct-link awareness	0.3921 ± 0.0330	0.1909	0.7417	90.38%	23.34%	32.01
Proposed reliability-aware DT	0.2076 ± 0.0042	0.0883	0.3052	90.38%	54.37%	12.00

4.2. Train-Validation-Test Protocol

The trace is divided into 60% training, 20% validation, and 20% test data. The first 8,640 samples, corresponding to three days, are used to fit the sensor-specific AR(2) models. The next 2,880 samples form the validation day. The final 2,880 samples form the held-out test day. This ordering prevents temporally adjacent future observations from being randomly mixed into model training.

The validation period is used to select γ and the AoI weight w_A. The chosen temporal-DT configuration is $\gamma = 0.10$ and $w_A = 0.40$. The direct-link-aware and proposed methods both use $\gamma = 0.25$ and $w_A = 0.80$. These values are frozen before the held-out test period.

4.3. Baselines and Wireless Delivery Model

Five methods are evaluated. Full direct transmission attempts every sensor update and uses last-value hold after packet loss. Periodic sampling uses round-robin scheduling within the same restricted radio budget and also uses last-value hold. The temporal DT uses the AR(2) predictor and AoI but no link information. The direct-link-aware DT adds $p_{i,0}$ to its update priority but does not relay. The proposed method includes both route reliability and radio cost and can select the cost-aware two-hop route from Eq. (4).

A scheduled direct packet is sampled as a Bernoulli event using the released $p_{i,0}$. A relayed update succeeds when both independently sampled hops succeed. The main result uses 30 matched seeds so that every method experiences the same set of random-number seeds. The communication-budget sweep uses 10 seeds per operating point. The use of aggregate probabilities means that the evaluation is trace-driven, not a replay of time-varying packet-level channel measurements.

5. Results and Discussion

5.1. Held-Out Reconstruction Accuracy

At $B = 5$, every restricted method uses 9.62% of the radio traffic required for full direct telemetry, equal to a 90.38% communication reduction. Periodic scheduling records 0.8176 ± 0.0677 RMSE. The temporal DT lowers RMSE to 0.4069 ± 0.0392 by predicting the state between successful physical updates. Direct-link awareness lowers the value further to 0.3921 ± 0.0330. The proposed method reaches 0.2076 ± 0.0042 under the same budget as shown in Table 4.

The direct-link-aware DT is the strongest ablation because it uses the same predictor family, AoI concept, and fixed radio budget. Relative to this method, the proposed route and cost logic reduces RMSE by 47.07% and reduces MAE from 0.1909 to 0.0883. This comparison isolates the practical value of recovering poorly connected nodes rather than attributing the full gain to the predictor. The full-direct baseline has 0.2764 mean RMSE even though every node attempts a transmission. Figure 2 shows the recorded lossy direct links. The proposed method obtain a lower network-level error because it spends two transmissions on selected weak nodes when the relay route provides better expected success per transmission. This does not mean that selective transmission contains more information than ideal full telemetry. A lossless full-telemetry oracle has zero reconstruction error by definition.

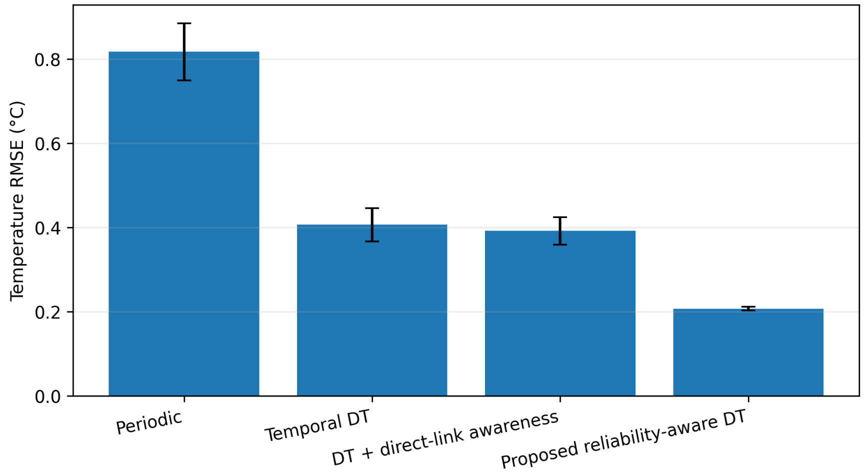

Figure 2: *Held-out temperature reconstruction RMSE at a radio budget of five transmissions per 30-second slot.*

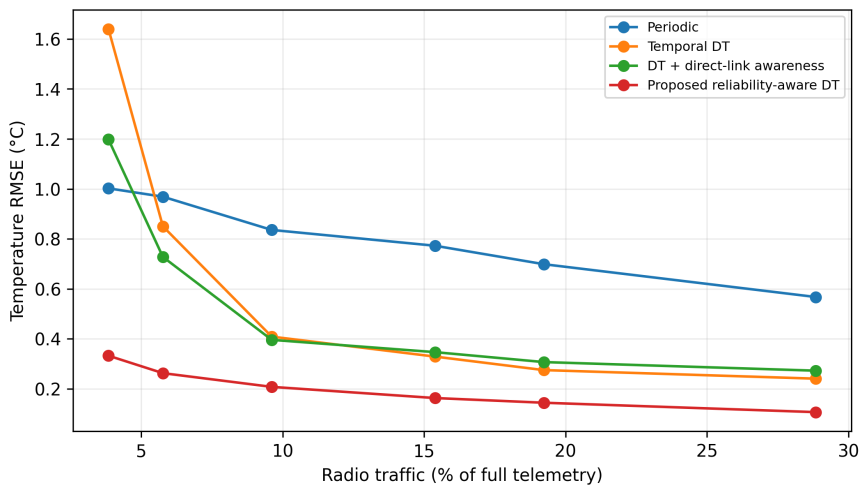

Figure 3: *Communication-accuracy trade-off across the tested radio budgets.*

5.2. Communication-Accuracy Trade-off

The proposed method remains the lowest-RMSE restricted approach across the evaluated 3.85%-28.85% traffic range. At two transmissions per slot, it uses only 3.85% of full traffic and obtains 0.3328 ± 0.0099 RMSE. At five transmissions, RMSE is approximately 0.208. Increasing the budget to 15 transmissions per slot reduces RMSE to 0.1071 ± 0.0035 while still removing 71.15% of full radio traffic. The curve in Figure 3 therefore exposes a usable operating range instead of selecting one threshold and claiming a universal optimum.

5.3. Freshness of the Virtual State

Reconstruction error alone can hide stale physical information, particularly for slowly varying temperature signals. At the main operating point, the temporal DT has a mean AoI of 35.46 slots and the direct-link-aware DT has 32.01 slots as shown in Figure 4. The proposed method reduces mean AoI to 12.00 slots. This is a 62.52% reduction relative to the strongest direct-link ablation. The 95th-percentile AoI also decreases from 74.67 to 28.87 slots.

The freshness gain is consistent with the route-selection logic. Direct-link awareness helps the scheduler avoid some poor update choices, but a sensor with a very weak direct path can still accumulate age. The proposed method can instead spend two transmissions on a selected relay when the expected reliability per transmission is higher. Because the radio budget is fixed, the gain comes from reallocating the same budget rather than increasing the number of allowed transmissions.

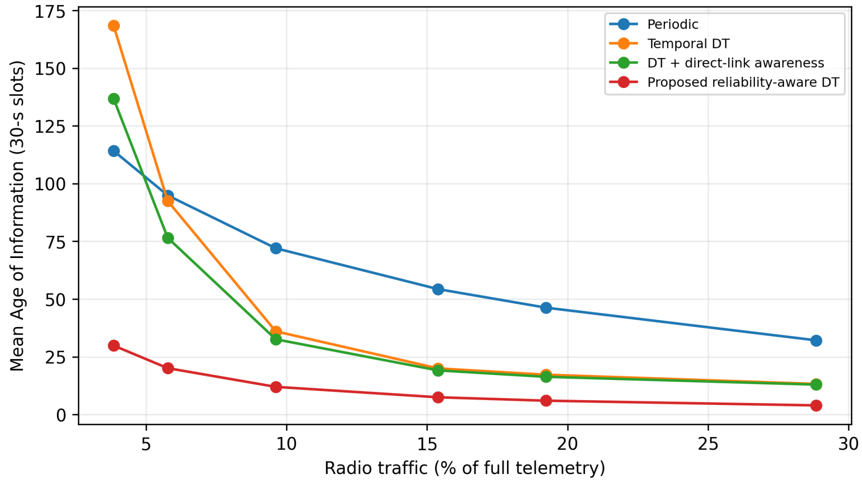

Figure 4: *Mean AoI versus radio traffic.*

Table 5. *Node-level effect of cost-aware relaying for the five nodes that select a two-hop route.*

node	Direct p	Relay	2-hop	Direct RMSE	Prop. RMSE	Direct AoI	Prop. AoI
14	0.0054	12	0.2432	0.9589	0.1538	170.6	26.4
16	0.0096	24	0.2259	0.9288	0.1666	104.8	12.5
17	0.0304	21	0.2865	0.4375	0.1483	38.7	11.6
18	0.0048	24	0.5058	1.2906	0.1109	200.3	9.5
19	0.0902	23	0.5478	0.2678	0.1415	35.8	17.2

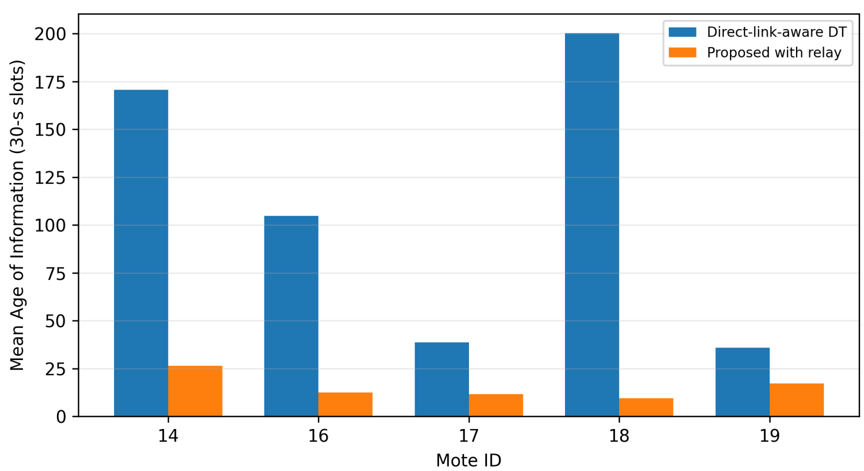

Figure 5: *Mean AoI of the five weak-link nodes under the direct-link-aware DT and the proposed cost-aware relay mechanism.*

5.4. Weak-Link Relay Ablation

Node 18 is the clearest weak-link case. Its recorded direct success probability is 0.0048 as shown in Table 5. Under direct-link-aware scheduling, node 18 records 1.2906 RMSE and 200.35-slot mean AoI. With the proposed path, these values decrease to 0.1109 and 9.50 slots. nodes 14, 16, 17, and 19 show the same direction of change, although the size of the improvement depends on their direct and relay probabilities.

Only five of the 52 active nodes satisfy the reliability-per-cost condition in Eq. (4). The remaining nodes continue to use direct delivery. Figure 5 result is useful for interpreting the method: the performance improvement is not produced by turning the network into a general multi-hop routing system. It comes from targeted recovery of a small number of weak direct links that would otherwise account for a large share of the state age and reconstruction error.

5.6. Discussion and Limitations

The results lead to three observations within the evaluated Intel trace.

- A temporal DT remains useful when only a small radio budget is available because it maintains an estimated sensor state between successful physical updates. The comparison with periodic scheduling shows that prediction reduces the error introduced by missed or suppressed measurements.
- Link reliability affects the value of a scheduled update. A sensor may have high priority because of its prediction behavior or AoI, but the update contributes little if its direct delivery probability is close to zero. The weak-link results show that accounting for the delivery path can substantially reduce this problem.
- Reconstruction error and state freshness provide complementary information. Under the same communication budget, the proposed method reduces both reconstruction error and AoI. Reporting these measures together gives a clearer view of synchronization quality than communication reduction alone.

This behavior also distinguishes the proposed method from recent adaptive sensing approaches [29]-[31]. Predictive suppression primarily decides whether a new physical measurement contains enough information to justify transmission. The proposed system makes an additional decision once an update is considered useful: it determines whether the direct or two-hop path provides sufficient delivery reliability for its radio cost. This distinction is relevant to the Intel topology because the five selected weak nodes have very low direct success probabilities. Repeated direct attempts at these nodes can therefore leave the corresponding DT states without a recent physical correction for long periods.

The evaluation has several limitations that define the scope of these results.

- The current paper evaluates the temperature channel only, although the original Intel deployment also recorded humidity, light, and voltage [32]. The reported results therefore should not be interpreted as multimodal DT synchronization performance.
- Packet delivery is sampled from aggregate link probabilities. The evaluation consequently does not reproduce time-correlated fading, interference, MAC-layer retransmissions, or contention among simultaneous wireless links.
- The two hops of a relay transmission are sampled independently. Receiver identifier 0 is modeled as the collection endpoint, as assumed in the experimental model.
- The Intel trace represents a well-established indoor wireless sensor network rather than a deployment based on contemporary Wi-Fi, LoRaWAN, 5G, or 6G radios. The measured gains therefore describe the considered trace-driven setting and should not be generalized directly to those radio technologies.

The paper therefore provides evidence for the proposed synchronization logic under the evaluated trace and connectivity model. It does not claim deployment-ready performance across all IoT networks.

6. Conclusion and Future Work

This paper presented a reliability-aware DT synchronization method for resource-constrained IoT sensor networks. The method combines an AR(2) virtual-state predictor with AoI, measured link reliability, and radio-transmission cost. Direct delivery is retained for most sensors. A two-hop relay is selected only when its expected success probability per radio transmission exceeds the direct alternative. The design therefore targets weak links without converting the full network to multi-hop forwarding.

The method was evaluated with a real five-day Intel Berkeley temperature trace and the corresponding aggregate connectivity measurements. At five transmissions per 30-second slot, equal to 9.62% of full telemetry, the proposed method obtains 0.2076 ± 0.0042 RMSE. The direct-link-aware DT obtains 0.3921 ± 0.0330 under the same budget, giving a 47.07% relative RMSE reduction. Mean AoI decreases from 32.01 to 12.00 slots, and the proposed approach remains the lowest-RMSE restricted method across the tested 3.85%-28.85% traffic range.

The current work is limited to one temperature trace and an aggregate-probability channel model. Future work will extend the evaluation to multiple sensing modalities, time-varying packet traces, and an embedded or emulated network. These experiments are needed before making claims about other radio technologies or deployment environments.

Author Declarations

Author Contributions (CRediT): Zupash Awais: Conceptualization, Methodology, Software, Validation, Formal Analysis, Visualization. Nauman Irshad Ali Shah: Writing - Original Draft, Writing - Review & Editing. All authors have read and approved the submitted version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The Intel Berkeley Research Lab sensor measurements, node locations, and connectivity data are publicly available from the Intel Lab Data repository [32].

Acknowledgments: The authors would like to thank the Center for Emerging Networks & Technologies (CENT) Lab at the University of Central Punjab (UCP), Lahore, Pakistan, for technical support and constructive feedback.

Conflicts of Interest: The authors declare no conflict of interest.

Ethics Statement: Not applicable.

Declaration of Generative AI: During the preparation of this manuscript, the authors used ChatGPT for language improvement. The authors reviewed and verified the resulting content and take full responsibility for the accuracy, integrity, and content of the manuscript.

References

[1] M. Bhatia and R. Kumar, "Digital Twin Technology in Wireless Communication: A Comprehensive Engineering Survey," IEEE Internet of Things Journal, vol. 12, no. 12, pp. 19144-19166, 2025, doi: 10.1109/JIOT.2025.3565843.

[2] D.-H. Tran et al., "Network Digital Twin for 6G and Beyond: An End-to-End View Across Multi-Domain Network Ecosystems," IEEE Open Journal of the Communications Society, vol. 6, pp. 6866-6911, 2025, doi: 10.1109/OJCOMS.2025.3599866.

[3] T. Li, Q. Long, H. Chai, S. Zhang, F. Jiang, H. Liu, W. Huang, D. Jin, and Y. Li, "Generative AI Empowered Network Digital Twins: Architecture, Technologies, and Applications," ACM Computing Surveys, vol. 57, no. 6, Art. no. 157, pp. 1-43, 2025, doi: 10.1145/3711682.

[4] Y. Sanjalawe, S. Fraihat, S. Al-E'mari, and S. N. Makhadmeh, "Bridging the gap: A comprehensive survey on AI-driven digital twin networks for future wireless systems," Journal of King Saud University Computer and Information Sciences, vol. 38, Art. no. 113, 2026, doi: 10.1007/s44443-026-00522-y.

[5] H. Tong, M. Chen, J. Zhao, Y. Hu, Z. Yang, Y. Liu, and C. Yin, "Continual Reinforcement Learning for Digital Twin Synchronization Optimization," IEEE Transactions on Mobile Computing, vol. 24, no. 8, pp. 6843-6857, 2025, doi: 10.1109/TMC.2025.3546507.

[6] Q. Zhang, Y. Wang, and Z. Li, "Scheduling of Digital Twin Synchronization in Industrial Internet of Things: A Hybrid Inverse Reinforcement Learning Approach," IEEE Internet of Things Journal, vol. 12, no. 5, pp. 5137-5147, 2025, doi: 10.1109/JIOT.2024.3486125.

[7] B. Li, H. Cai, L. Liu, and Z. Fei, "Delay-Aware Digital Twin Synchronization in Mobile Edge Networks With Semantic Communications," IEEE Transactions on Vehicular Technology, vol. 74, no. 7, pp. 10974-10983, 2025, doi: 10.1109/TVT.2025.3548844.

[8] H. Yu, Y. Liu, Z. Yang, H. Sun, and M. Chen, "Optimizing Wireless Resource Management and Synchronization in Digital Twin Networks," IEEE Internet of Things Journal, vol. 12, no. 15, pp. 29152-29163, 2025, doi: 10.1109/JIOT.2025.3543126.

[9] Y. Li, W. Liang, Z. Xu, W. Xu, and X. Jia, "Budget-Constrained Digital Twin Synchronization and Its Application on Fidelity-Aware Queries in Edge Computing," IEEE Transactions on Mobile Computing, vol. 24, no. 1, pp. 165-182, 2025, doi: 10.1109/TMC.2024.3455357.

[10] A. Loubany, M. Itani, and S. Sharafeddine, "From Age of Information to Age of Digital Twin: A Review on Synchronization Metrics for IoT Networks," IEEE Access, vol. 13, pp. 131102-131119, 2025, doi: 10.1109/ACCESS.2025.3591589.

[11] K. Noroozi, T. D. Todd, D. Zhao, and G. Karakostas, "Age of Information in Digital Twin Migration," IEEE Transactions on Vehicular Technology, vol. 74, no. 10, pp. 16281-16294, 2025, doi: 10.1109/TVT.2025.3570505.

[12] Y. Xiao, Q. Du, and G. K. Karagiannidis, "Statistical Age of Information: A Risk-Aware Metric and Its Applications in Status Updates," IEEE Transactions on Wireless Communications, vol. 24, no. 3, pp. 2325-2340, 2025, doi: 10.1109/TWC.2024.3520324.

[13] Y. Wang, L. Tang, L. Wang, H. Zhang, W. Wang, D. Fang, and Q. Chen, "Digital Twin Information Synchronization Strategy for IIoT Based on Dual-Time-Scale Network Slicing Orchestration," IEEE Internet of Things Journal, vol. 13, no. 2, pp. 1919-1935, 2026, doi: 10.1109/JIOT.2025.3625762.

[14] J. Li, J. Wang, W. Liang, Q. Chen, S. K. Das, and X. Jia, "Digital Twin Freshness Maximization in Edge Computing," IEEE Transactions on Services Computing, vol. 19, no. 2, pp. 1107-1119, 2026, doi: 10.1109/TSC.2026.3651602.

[15] G. Khalaf, M. Itani, and S. Sharafeddine, "A UAV-Aided Digital Twin Framework for IoT Networks With High Accuracy and Synchronization," IEEE Transactions on Network and Service Management, vol. 23, pp. 3013-3025, 2026, doi: 10.1109/TNSM.2026.3670040.

[16] Y. Zhang, L. Wang, and W. Liang, "Deep Reinforcement Learning for Mobility-Aware Digital Twin Migrations in Edge Computing," IEEE Transactions on Services Computing, vol. 18, no. 2, pp. 704-717, 2025, doi: 10.1109/TSC.2025.3528331.

[17] X. Ai, W. Liang, Y. Zhang, and W. Xu, "Fidelity-Aware Inference Services in DT-Assisted Edge Computing via Service Model Retraining," IEEE Transactions on Services Computing, vol. 18, no. 4, pp. 2089-2102, 2025, doi: 10.1109/TSC.2025.3586126.

[18] J. Li, J. Wang, W. Liang, X. Jia, and A. Y. Zomaya, "Inference Service Fidelity Maximization in DT-Assisted Edge Computing," IEEE Transactions on Mobile Computing, vol. 25, no. 1, pp. 1352-1366, 2026, doi: 10.1109/TMC.2025.3600390.

[19] X. Ai, W. Liang, and C. Liu, "Joint Optimization of Model Retraining and Inference Services in DT-Assisted Edge Computing," IEEE Transactions on Networking, vol. 34, pp. 1804-1819, 2026, doi: 10.1109/TON.2025.3632228.

[20] Y. Shi, L. Lao, X. Zheng, et al., "Real-time Spatiotemporal Synchronization for Human-robot Collaborative Systems: An Integrated Digital Twin Framework with Adaptive Data Processing," Journal of Intelligent & Robotic Systems, vol. 112, Art. no. 12, 2026, doi: 10.1007/s10846-025-02346-w.

[21] A. Ndikumana, K.-K. Nguyen, and M. Cheriet, “Digital Twin Backed Closed-Loops for Energy-Aware and Open RAN-Based Fixed Wireless Access Serving Rural Areas,” IEEE Transactions on Mobile Computing, vol. 24, no. 3, pp. 1669-1683, 2025, doi: 10.1109/TMC.2024.3482985.

[22] M. A. Qadir, M. Naeem, and W. Ejaz, “Digital twin-assisted multi-layer networks for low-latency and energy-efficient communication,” Computer Communications, vol. 241, Art. no. 108219, 2025, doi: 10.1016/j.comcom.2025.108219.

[23] S. Chu, J. Li, J. Wang, Y. Ni, K. Wei, W. Chen, and S. Jin, “Resource Efficient Asynchronous Federated Learning for Digital Twin Empowered IoT Network,” IEEE Transactions on Green Communications and Networking, vol. 10, pp. 201-217, 2026, doi: 10.1109/TGCN.2025.3574143.

[24] Y. Zhou, Y. Fu, Z. Shi, H. H. Yang, K. Hung, and Y. Zhang, “Energy-Efficient Federated Learning and Migration in Digital Twin Edge Networks,” IEEE Transactions on Cognitive Communications and Networking, vol. 12, pp. 743-758, 2026, doi: 10.1109/TCCN.2025.3567096.

[25] Y. Kwon, “Smart Manufacturing Innovation Through the Integration of Digital Twin and IoT Energy Management,” International Journal of Control, Automation, and Systems, vol. 23, no. 11, pp. 3427–3439, 2025, doi: 10.1007/s12555-025-0046-1.

[26] L. Chang, X. Yang, K. Zhang, et al., “Blockchain-based synchronized data transmission with dynamic device management for digital twin in IIoT,” Peer-to-Peer Networking and Applications, vol. 18, Art. no. 85, 2025, doi: 10.1007/s12083-025-01910-3.

[27] M. Yahya, M. Naeem, R. H. Alsisi, M. O. Butt, and W. Ejaz, “Robust Task Offloading in Digital Twin-Enabled UAV-IoT Networks Under Spatial Uncertainty,” IEEE Open Journal of the Communications Society, vol. 7, pp. 1717–1728, 2026, doi: 10.1109/OJCOMS.2026.3663770.

[28] K. Hayawi, J. Sajid, A. W. Malik, and S. S. Mathew, “Digital Twin-Assisted Task Offloading for Workload Management at Fog Nodes,” IEEE Internet of Things Journal, vol. 12, no. 13, pp. 23061-23072, 2025, doi: 10.1109/JIOT.2025.3550832.

[29] F. Alanazi and M. N. Alanazi, “Machine learning driven aggregation aware bitmap MAC protocol for energy efficient data transmission in WSNs,” Scientific Reports, vol. 15, Art. no. 36893, 2025, doi: 10.1038/s41598-025-20835-8.

[30] G. Lu, Y. Sun, Y. Jiang, and B. Butler, “Adaptive Sampling for Spatiotemporal Anomaly Monitoring in Wireless Sensor Networks,” arXiv:2607.15235, 2026. Accepted for IEEE ISSC 2026.

[31] J. Kangethe, I. I. Uddin, and L. Wang, “Learning to Transmit: Volatility-Aware Predictive Communication for Energy-Efficient IoT Networks,” arXiv:2607.19590, 2026.

[32] Intel Berkeley Research Lab, “Intel Lab Data,” MIT Computer Science and Artificial Intelligence Laboratory. Accessed: Aug. 14, 2026. [Online]. Available: https://db.csail.mit.edu/labdata/labdata.html.