Global Journal of Engineering and Technology Advances
eISSN: 2582-5003
Cross Ref DOI: 10.30574/gjeta
Journal homepage: https://gjeta.com/

(RESEARCH ARTICLE)

Check for updates

MEASUREMENT-UNCERTAINTY-DRIVEN TRANSMISSION SCHEDULING FOR INDUSTRIAL WIRELESS SENSOR NETWORKS: A STABILITY DICHOTOMY IN FRESHNESS-FAIR ALLOCATION

Md. Masum Hasan Maruf [1,*], Md. Mamun Ur Rashid [2], Md. Thouhedul Islam Rafi [3], Shah Syeda Sawaira [3] and Tasmia Aisha Siddika [2]

[1] School of Electronics and Information Engineering, Nanjing University of Information Science and Technology, 219 Ningliu Road, Nanjing 210044, Jiangsu, China.
[2] School of Computer Science and Technology, Nanjing University of Information Science and Technology, 219 Ningliu Road, Nanjing 210044, Jiangsu, China.
[3] Department of Computer Science and Technology, School of Computer Science, Jiangsu University of Science and Technology, 666 Changhui Road, Zhenjiang 212100, Jiangsu, China.
* Corresponding Author

ORCID Details
Md. Masum Hasan Maruf: https://orcid.org/0009-0009-0374-2716
Md. Mamun Ur Rashid: https://orcid.org/0009-0000-8836-0673
Md. Thouhedul Islam Rafi: https://orcid.org/0009-0008-5765-3869
Tasmia Aisha Siddika: https://orcid.org/0009-0000-7443-3368

Global Journal of Engineering and Technology Advances, 2026, 29(02), 119–131

Publication history: Received on 11 October 2026; revised on 21 November 2026; accepted on 23 November 2026

Article DOI: https://doi.org/10.30574/gjeta.2026.29.2.0333

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

Age of Information (AoI), widely used for scheduling in industrial wireless sensor networks, reflects when a sample was generated but not the accuracy of the instrument that produced it, although instruments on one network can differ in accuracy by orders of magnitude and degrade with age. We determine when freshness-fair scheduling is optimal for remote state estimation and how the allocation should change when it is not. A time-slotted channel serving scalar Gauss–Markov sources observed through instruments of different noise variance is modelled, the periodic estimation cost is derived in closed form, and a continuous-rate relaxation of the allocation problem is solved by bisection on the shadow price of a slot. For integrating sources with common dynamics, freshness-fair allocation is proved strictly suboptimal whenever instrument variances differ; for random-walk sources the converse is also proved, and the optimal rate is shown to increase with instrument noise. For self-regulating sources, the marginal value of a slot is proved to be bounded by an instrument information budget that decreases with instrument noise; under a numerically verified, unproved convexity property, this yields an abandonment threshold that is explicit once the slot price is known. In synthetic tests, a water-filling policy reduces total mean-square error by 11.5% relative to round robin and is within 0.5% of a numerically computed Whittle-index scheduler; with measured 36-month uncertainty trajectories from a gas-sensor array and modelled dynamics, the reduction is 3.5%. Instrument uncertainty can therefore be a useful input to industrial wireless scheduling.

Keywords: Age of information; Industrial wireless sensor networks; Measurement uncertainty; Sensor scheduling; Remote state estimation; Time-slotted channel hopping

1 INTRODUCTION

Low-power wireless networks increasingly carry measurement traffic in process plants. WirelessHART (IEC 62591), ISA100.11a, WIA-PA and the IETF 6TiSCH stack use time-slotted access with channel hopping [1,2], in which a central network manager assigns each device a repeating pattern of transmission slots. Deciding which sensor transmits in which slot is therefore a central resource-allocation problem in an industrial wireless sensor network (IWSN).

Age of Information (AoI) [3,4], the time elapsed since the freshest delivered update was generated, is a widely used metric for this decision; it need not be minimised by zero-wait updating [5] and has been optimised under broadcast, multi-source, energy-harvesting and finite-blocklength constraints [6–9]. AoI, however, records when a sample was taken, not how accurate it is. Instruments on the same network may differ in measurement variance by orders of magnitude (Table 1), and an instrument's uncertainty grows as its calibration ages. For a source with prior error covariance P^- observed through an instrument of noise variance R, a Kalman update reduces the covariance by

$$\Delta = P^- - \frac{P^- R}{P^- + R} = \frac{(P^-)^2}{P^- + R}. \tag{1}$$

Staleness enters only through P^- and the instrument only through R; as $R \to \infty$ the reduction vanishes regardless of age. It is known that AoI-optimal sampling need not minimise estimation error, because a sampler can exploit knowledge of the signal [10,11]. Here we study a different source of mismatch that remains when the scheduler is signal-agnostic: heterogeneity of the instruments.

1.1 Related work and gap

Beyond the scheduling results surveyed in [9], including Whittle-index treatments of age functions [8] and industrial variants with bounded peak age [12], tail constraints [13] and duty-cycling [14], several metrics replace age by a measure of relevance: Age of Incorrect Information [15], Urgency of Information [16], Uncertainty of Information [17], Value of Information [18–20] and goal-oriented communication [21,22]. In these metrics the uncertainty resides at the receiver or in the application. The Uncertainty of Information of [17], for example, is the monitor's entropy about the source, whereas the uncertainty considered here is the instrument's metrological uncertainty in the sense of the GUM [23]. In remote estimation, Kalman filtering with intermittent observations [24], estimation with communication costs [25], event-triggered scheduling [26,27], multi-sensor scheduling on shared channels [28] and Whittle-index multichannel allocation [29] are well established, and sensor selection already accounts for different noise variances [30]. In that literature, however, the noise variance is usually a fixed design constant. Metrology, conversely, provides in-situ uncertainty estimation [31], machine-readable calibration certificates [32] and condition-based recalibration [33], and quality-of-information frameworks anticipated related concerns [34], but this information is rarely used for channel allocation. The gap we address is narrow: how the long-run slot allocation of a shared channel should respond to differences, and drift, in instrument variance.

1.2 Objective

The objective of this study is to establish when freshness-fair scheduling is optimal for remote estimation over a shared industrial wireless channel, to characterise the optimal allocation when it is not, and to obtain a policy that realises it with integer slot grants. Our working hypothesis is that heterogeneity in instrument uncertainty, not staleness alone, should determine how bandwidth is divided, and that the direction of the correct response depends on the open-loop stability of the monitored process.

1.3 Contributions

Within the scalar linear-Gaussian model of Section 2, and for a continuous-rate relaxation of the slot-allocation problem:

For integrating sources ($|a| \geq 1$, with a the source pole) with common dynamics, we prove that freshness-fair allocation is strictly suboptimal whenever instrument variances differ, and, for random-walk sources ($a^2 = 1$), that it is optimal when they coincide (Theorem 1). A numerical example shows that the first statement does not extend to stable sources.

For random-walk sources we prove that the optimal rate increases strictly with instrument variance (Theorem 2). For stable sources we prove that the marginal value of a slot is bounded by an instrument information budget that tends to zero as the instrument degrades (Proposition 2), and, conditional on a numerically verified convexity property, we derive an abandonment threshold (Proposition 3).

We give a water-filling allocation computed by scalar bisection and realised as integer slot grants, and evaluate it against round robin, greedy rules and a numerically computed Whittle-index scheduler, including sensitivity to misreported variances, lossy links and measured sensor drift.

We distinguish throughout between proved results and numerical observations; Table 2 lists the status of each.

2 MATERIALS AND METHODS

2.1 System model

N devices report to a gateway over a shared slotted channel. On each of B parallel channels one device may transmit per slot; $f_i \in [0,1]$ is the long-run fraction of slots assigned to device i, so $\sum_i f_i = B$. Figure 1 shows the loop: the per-device measurement variance R_i comes from the instrumentation layer (a calibration certificate, accuracy class or online estimate) and is combined with the slot allocation at the gateway. R_i is assumed known in the analysis; the allocation is recomputed when an estimate changes materially.

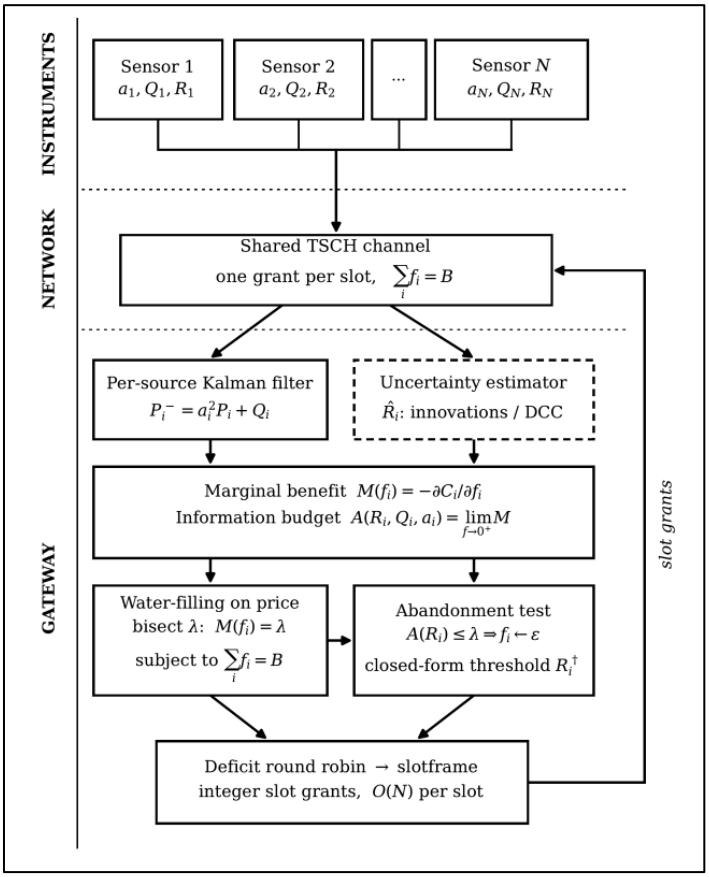

Figure 1: Three-layer architecture of the proposed scheduling loop.

Device i observes a scalar Gauss–Markov process through a noisy instrument,

$$x_i(k+1) = a_i x_i(k) + w_i(k), \qquad w_i \sim \mathcal{N}(0, Q_i), \tag{2}$$

$$y_i(k) = x_i(k) + v_i(k), \qquad v_i \sim \mathcal{N}(0, R_i), \tag{3}$$

where Q_i is the process-noise variance and R_i the measurement-noise (instrument) variance. The gateway runs a per-source Kalman filter that applies the prediction step every slot and the update step when device i is served:

$$P_i^-(k) = a_i^2 P_i(k-1) + Q_i, \tag{4}$$

$$P_i(k) = \begin{cases} \dfrac{P_i^-(k)\,R_i}{P_i^-(k) + R_i}, & \text{if served,} \\ P_i^-(k), & \text{otherwise.} \end{cases} \tag{5}$$

A self-regulating variable such as a regulated flow has $|a_i| < 1$ and bounded open-loop variance; an integrating variable such as a tank level has $|a_i| \geq 1$. All quantities below depend on a_i only through a_i^2, and we call $a_i^2 = 1$ the random-walk case.

Table 1 gives the range of R used in the synthetic study. Each value is obtained from a published accuracy class by treating the specification limit as a coverage interval and dividing by the coverage factor (GUM Type B [23]), normalised to a calibrated pressure transmitter under a shared span. The entries are order-of-magnitude scalings, not models of specific instruments.

Table 1: Accuracy class mapped to relative measurement-noise variance. RTD: resistance temperature detector; MEMS: micro-electro-mechanical system.

Instrument	Stated accuracy	Relative R
Pressure transmitter (calibrated)	±0.05% of span	1
Coriolis flow meter	±0.10% of rate	4
Pt100 RTD, Class A	±0.15 °C	$\sim 10^2$
Type K thermocouple, Class 1	±1.5 °C	$\sim 10^3$
Low-cost MEMS accelerometer	noise-density limited	$\sim 10^4$

2.2 Problem formulation

Let $C(f; R, Q, a)$ be the time-average error covariance of a source served at rate f (defined in Proposition 1). The allocation problem is

$$\min_f \sum_i C\,(f_i; R_i, Q_i, a_i) \quad \text{s.t.} \quad \sum_i f_i = B, \quad 0 \leq f_i \leq 1. \tag{6}$$

We call the objective the total mean-square error (MSE); the freshness-fair allocation is $f_i = B/N$. Problem (6) is a continuous-rate relaxation: C is exact for strictly periodic service with integer period $T = 1/f$, and for non-integer $1/f$ it is defined by the same expression, which interpolates between integer periods. All analytical statements below refer to (6); the gap to integer schedules is measured in Sections 2.4 and 3.4. We write λ for the Lagrange multiplier of the budget constraint, interpreted as the shadow price of a slot.

2.3 Analysis

We reserve Lemma, Proposition, Theorem and Corollary for statements with complete proofs. Statements supported only by numerical evidence are labelled Observation, and any result that uses one says so.

2.3.1 Periodic cost

Proposition 1. For a source served every T slots, let $\alpha = a^{2T}$, $G = \sum_{j=0}^{T-1} a^{2j} = (\alpha - 1)/(a^2 - 1)$ ($G = T$ if $a^2 = 1$) and $S = QG$. The periodic post-update covariance p_0 is the unique positive root of

$$\alpha p_0^2 + p_0(S + R - \alpha R) - SR = 0, \tag{7}$$

and the time-average covariance is

$$C(T) = \frac{1}{T}\left[p_0 G + \frac{Q(G - T)}{a^2 - 1}\right], \tag{8}$$

which reduces to $C(T) = p_0 + Q(T - 1)/2$ when $a^2 = 1$.

Proof. From p_0, (4) gives $P(j) = a^{2j}p_0 + Q\sum_{l<j} a^{2l}$, $j = 0, \ldots, T-1$, and the pre-update covariance after T steps is $u = \alpha p_0 + S$. Periodicity requires $p_0 = uR/(u+R)$, which rearranges to (7); since $\alpha > 0$ and $-SR < 0$, (7) has exactly one positive root. Averaging $P(j)$ and using $\sum_j \sum_{l<j} a^{2l} = (G-T)/(a^2-1)$ gives (8). ▫

Numerically, the root should be evaluated as $p_0 = 2SR/\left(b + \sqrt{b^2 + 4\alpha SR}\right)$ with $b = S + R - \alpha R$, because the textbook form suffers cancellation when α is small. Equation (8) agrees with direct Riccati iteration to a relative error below 10^{-6}. With

$$f = \frac{1}{T}, \tag{9}$$

defines $C(f) := C(1/f)$ and $\partial C/\partial f = -T^2\,\partial C/\partial T$; $T \geq 1$ is treated as continuous in all derivatives.

2.3.2 Marginal benefit of a slot

Let $M(f) = -\partial C/\partial f$ denote the marginal reduction in MSE per unit rate. For $a^2 = 1$,

$$M(f; R, Q) = \frac{\sqrt{Q}(Q/2 + Rf)}{f^2\sqrt{Q + 4Rf}}. \tag{10}$$

Lemma 1. For $|a| \geq 1$, M is strictly increasing in R at every fixed $f \in (0,1]$.

Proof. Let $\beta = a^2 \geq 1$; then $\alpha = \beta^T$ is non-decreasing and $S = QG$ strictly increasing in T. By Proposition 1, $u = \alpha p_0 + S$ is the positive root of $h(u) := u - \alpha uR/(u+R) - S$. (i) Multiplying $h(u) = 0$ by $u + R$ gives $u(u + R - \alpha R) = S(u+R) > 0$, so $u + R > \alpha R \geq \sqrt{\alpha}R$; hence $D := (u+R)^2 - \alpha R^2 > 0$ and $h_u = D/(u+R)^2 > 0$. (ii) Since $h_\alpha < 0$ and $h_S = -1$, the implicit function theorem gives $du/dT > 0$. (iii) Differentiating $p_0 = uR/(u+R)$ with $du/dR = \alpha\, dp_0/dR$ gives $\Psi := \partial p_0/\partial R = u^2/D > 0$, with $\partial\Psi/\partial u = 2uR(u + R - \alpha R)/D^2 > 0$ and $\partial\Psi/\partial\alpha = u^2R^2/D^2 > 0$, so $d\Psi/dT > 0$. (iv) In (8) only p_0 depends on R, so $\partial C/\partial R = (G/T)\Psi$. $G/T = 1$ if $\beta = 1$, and for $\beta > 1$ it is strictly increasing in T because $\frac{d}{dT}\frac{\beta^T - 1}{T} = \frac{e^x(x-1)+1}{T^2} > 0$ with $x = T\ln\beta$. Hence $\partial C/\partial R$ is strictly increasing in T, and $\partial M/\partial R = T^2\,\partial(\partial C/\partial R)/\partial T > 0$. ▫

Lemma 2. For $a^2 = 1$, M is strictly decreasing in $f > 0$, with $M \to \infty$ as $f \to 0^+$ and $M \to 0$ as $f \to \infty$.

Proof. Differentiating (10),

$$\frac{dM}{df} = -\frac{\sqrt{Q}(Q^2 + 6QRf + 6R^2f^2)}{f^3\,(Q + 4Rf)^{3/2}} < 0. \tag{11}$$

The limits follow from (10). ▫

Corollary 1. If $a_i^2 = 1$ for all i, (6) is strictly convex; its Karush–Kuhn–Tucker (KKT) conditions are necessary and sufficient and its minimiser is unique.

Proof. $dM/df < 0$ is equivalent to $\partial^2 C/\partial f^2 > 0$; a separable sum of strictly convex functions is strictly convex and the constraints are linear. ▫

Observation 1 (numerical). For $a^2 \neq 1$, M is non-increasing in f (equivalently, C is convex in f), and for $|a| < 1$, $M(f) < A$ for all $f > 0$, with A as in Proposition 2. This was checked on 37,680 comparisons over $a \in [0.3, 1.4]$, $R \in [10^{-6}, 10^6]$, $Q \in [0.01, 100]$ and $f \in [10^{-3}, 1]$ (Section 2.5). We have no proof.

2.3.3 When is freshness-fair allocation optimal?

Theorem 1. Let $a_i = a$, $Q_i = Q$ for all i, and $N > B$. (i) If $a^2 = 1$ and all R_i are equal, $f_i = B/N$ is the unique optimum of (6). (ii) If $|a| \geq 1$ and $R_i \neq R_j$ for some pair, $f_i = B/N$ is strictly suboptimal.

Proof. (i) The problem is permutation-symmetric and, by Corollary 1, has a unique minimiser, which must therefore be symmetric. (ii) At the uniform allocation all rates equal $B/N \in (0,1)$. If $R_i > R_j$, Lemma 1 gives $M(B/N; R_i) > M\left(B/N; R_j\right)$, so moving a small rate η from j to i changes the objective by $-\eta\left[M(B/N; R_i) - M\left(B/N; R_j\right)\right] + o(\eta) < 0$. ▫

Part (i) extends to other a if Observation 1 holds; part (ii) needs no convexity. Part (ii) does not extend to stable sources. Example: for $a = 0.95$, $Q = 1$ and $f = 0.5$, $M = 1.92289$ both at $R = 0.1$ and at $R \approx 354.37$ (computed numerically). A two-device single-channel network with these instruments satisfies the first-order condition at $f_1 = f_2 = 0.5$, and under Observation 1 this uniform allocation is optimal although the variances differ by more than three orders of magnitude.

2.3.4 Stable sources: the instrument information budget

For $|a| < 1$ let

$$\sigma_{ol}^2 = \frac{Q}{1 - a^2} \tag{12}$$

denote the open-loop variance, i.e. the error covariance of an estimator that receives no measurements.

Proposition 2. Let $|a| < 1$. Then C extends continuously to $f = 0$ with $C(0) = \sigma_{ol}^2$, and

$$\lim_{f \to 0^+} M(f) = A(R, Q, a) := \frac{(\sigma_{ol}^2)^2}{(1 - a^2)(\sigma_{ol}^2 + R)}, \tag{13}$$

with $|M(f) - A| = O(Ta^{2T})$. A is strictly decreasing in R and tends to 0 as $R \to \infty$.

Proof. Using $Q/(a^2 - 1) = -\sigma_{ol}^2$, (8) becomes $C(f) = \sigma_{ol}^2 - f\,\Phi(1/f)$ with $\Phi(T) := G(\sigma_{ol}^2 - p_0)$. As $T \to \infty$, $\alpha \to 0$, $G \to 1/(1 - a^2)$, $u \to \sigma_{ol}^2$ and $p_0 \to \sigma_{ol}^2 R/(\sigma_{ol}^2 + R)$, so $\Phi \to A$ and, Φ being bounded, $C(0^+) = \sigma_{ol}^2$. Moreover $M = \Phi(T) - T\Phi'(T)$; every T-dependence of Φ enters through $\alpha = a^{2T}$, so $\Phi' = O(a^{2T})$ and $\Phi - A = O(a^{2T})$. The properties of A follow from (13). ▫

We call A the instrument information budget: for small f, $C(0) - C(f) \approx Af$, so A is the MSE reduction per unit of a vanishing service rate. It is finite only for stable sources and has the same units as λ. Two further numerical properties are visible in Figure 2.

Observation 2 (numerical). For $|a| < 1$, M is not monotone in R in general: at any fixed f, $M \to 0$ as $R \to \infty$, but above a rate $f_c(a)$ it first rises with R over an intermediate range. On $R \in [10^{-3}, 10^4]$ with $Q = 1$ we obtain $f_c \approx 0.54$, 0.35, 0.17 and 0.033 for $a = 0.85$, 0.90, 0.95 and 0.99.

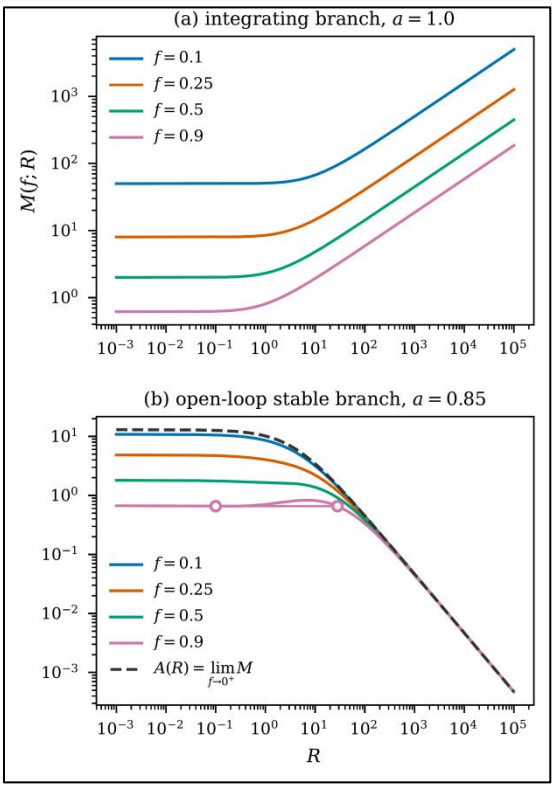

Figure 2: Marginal benefit $M(f; R)$ against instrument variance: (a) $a = 1$; (b) $a = 0.85$, with $A(R)$ dashed. In (b), the two marked points at $f = 0.9$ have equal M at variances about 283 times apart.

2.3.5 How the optimal rate responds to instrument variance

Theorem 2 (random-walk sources). Let $a_i^2 = 1$. (a) For every $\lambda > 0$, $M(f; R_i, Q_i) = \lambda$ has a unique solution $f_i(\lambda) > 0$, strictly increasing in R_i and strictly decreasing in λ; the rate used in (6) is $\min\{f_i(\lambda), 1\}$, which is positive, so such a source is never abandoned. (b) If $a_j^2 = 1$ for all j and no rate is at the cap at the optimum (e.g. $B = 1, N \geq 2$), increasing R_i strictly increases f_i^* and strictly decreases every other f_j^*.

Proof. (a) By Lemma 2, M decreases strictly and continuously from ∞ to 0, so the solution exists and is unique; the implicit function theorem with Lemmas 1 and 2 gives the signs. (b) By Corollary 1 the optimum satisfies $\sum_j f_j(\lambda^*) = B$. If R_i increases and the new price satisfied $\lambda' \leq \lambda^*$, every rate would weakly increase and f_i strictly, violating the budget; hence $\lambda' > \lambda^*$, every $f_j, j \neq i$, strictly decreases, and f_i must increase. ▫

Proposition 3 (stable sources; conditional on Observation 1). Assume Observation 1 holds for every device and let λ^* be the optimal multiplier in (6). A device with $|a_i| < 1$ has $f_i^* = 0$ if and only if $A(R_i, Q_i, a_i) \leq \lambda^*$, i.e. if and only if $R_i \geq R^\dagger$ with

$$R^\dagger = \frac{(\sigma_{ol}^2)^2}{\lambda^*(1 - a^2)} - \sigma_{ol}^2. \tag{14}$$

Proof. Under Observation 1 each C_i is convex, so the KKT conditions are necessary and sufficient. By Proposition 2 the right-derivative of C_i at 0 is $-A_i$, so $f_i = 0$ satisfies them iff $A_i \leq \lambda^*$. Solving $A(R^\dagger) = \lambda^*$ gives (14). ▫

Theorem 2 and Proposition 3 point in opposite directions. For a random-walk source, a noisier instrument leaves a larger residual covariance after each update, so additional slots are worth more. For a stable source, a sufficiently noisy measurement adds little beyond what the process model already provides, so the slot is better used elsewhere. Note that $R^\dagger$ is explicit only given λ^*, which depends on all devices and is found by bisection; it is a network quantity, not a per-sensor constant. Table 2 summarises the status of all analytical results.

Table 2: Status of the analytical results

Result	Scope	Status
Proposition 1 (periodic cost)	all a	proved
Lemma 1 (M increasing in R)	$\lvert a\rvert \geq 1$	proved
Lemma 2, Corollary 1 (convexity)	$a^2 = 1$	proved
Observation 1 (convexity)	$a^2 \neq 1$	numerical
Theorem 1(i)	$a^2 = 1$ (other a: under Obs. 1)	proved
Theorem 1(ii)	$\lvert a\rvert \geq 1$	proved
Proposition 2 (information budget)	$\lvert a\rvert < 1$	proved
Observation 2	$\lvert a\rvert < 1$	numerical
Theorem 2	$a^2 = 1$	proved
Proposition 3 (abandonment)	$\lvert a\rvert < 1$	conditional on Obs. 1

2.4 The MUD scheduling policy

The measurement-uncertainty-driven (MUD) policy equalises marginal benefit across devices. Because M is decreasing in f (Lemma 2 for $a^2 = 1$; Observation 1 otherwise), the rates of (6) can be found by bisection on λ (Algorithm 1). A streaming variant serves $\operatorname{argmax}_i I_i(k)$ with

$$I_i(k) = M\big(\hat{f}_i(k); R_i, Q_i, a_i\big), \qquad \hat{f}_i(k) = \frac{s_i(k)}{k}, \tag{15}$$

where $s_i(k)$ is the number of slots granted to device i up to slot k.

Algorithm 1. MUD slot allocation. Input: R_i, Q_i, a_i; budget B; probing floor ϵ; tolerance δ.

Bracket λ: $\lambda_{lo} \leftarrow 0$, $\lambda_{hi} \leftarrow 1$; double λ_{hi} while $\sum_i f_i(\lambda_{hi}) > B$.

Repeat: $\lambda \leftarrow \sqrt{\lambda_{lo}\lambda_{hi}}$; for each i, solve $M(f; R_i, Q_i, a_i) = \lambda$, cap at 1 and set $f_i \leftarrow \max(\epsilon, f)$; update the bracket by the sign of $\sum_i f_i - B$; until the bracket is within δ.

Set $f_i \leftarrow \epsilon$ for every i with $A(R_i, Q_i, a_i) \le \lambda$, and rescale the remaining rates so that $\sum_i f_i = B$.

For each slot k: $d_i \leftarrow d_i + f_i$ for all i; grant the slot to $j = \mathrm{argmax}_i d_i$; $d_j \leftarrow d_j - 1$ (deficit round robin, d_i initialised to 0).

Recompute when an estimate R_i changes materially.

Slot realisation and cost. Step 4 yields integer slot counts in a repeating frame. For the second configuration of Table 3, a 32-slot frame receives (6,6,5,15) slots versus (8,8,8,8) for round robin and is within 1.6% of the continuous-rate optimum of (6). Each allocation needs $O\big(N\log(1/\delta)\big)$ evaluations of M; the per-slot step is $O(N)$. In an unoptimised Python implementation on a general-purpose CPU, the per-slot grant took 1.74–1.87 μs for $N = 4-64$ and a full reallocation 37–640 ms. We have not implemented the policy on a TSCH stack; we note only that its output is a static slot assignment of the kind a network manager distributes.

Probing floor. With $f_i = 0$ an abandoned sensor's age grows without limit and, because R_i is estimated from innovations, the scheduler could not detect a recalibration. Every device is therefore guaranteed a minimum rate ϵ.

Lossy links. With packet error rate p_i and success probability $q_i = 1 - p_i$, lost updates being discarded, we approximate the cost by that of lossless service at the effective rate $f_i q_i$. Stationarity of $\sum_i C(f_i q_i)$ gives

$$q_i\, M(f_i q_i) = \lambda, \tag{16}$$

so the same bisection applies at price λ/q_i. This effective-rate cost is a surrogate; its error against simulated Bernoulli losses is reported in Section 3.5.

2.5 Computational procedures and data

All methods were implemented in Python. Unless stated otherwise, $N = 4$, $Q_i = 1$ and $B = 1$. The baselines are round robin; greedy AoI (serve the stalest source); greedy covariance (serve $\mathrm{argmax}_i P_i^-$); greedy Δ (serve $\mathrm{argmax}_i \Delta_i$ of (1)); and a Whittle-index scheduler. For the Whittle index, the single-arm average-cost Bellman equation with passivity subsidy m,

$$h(P) + g = \min\big\{c_a(P) + h(f_a(P)),\ c_p(P) - m + h(f_p(P))\big\}, \tag{17}$$

where f_a, f_p are the served and idle maps of (4)–(5) and the per-slot cost is the resulting covariance, was solved by relative value iteration on a 600-point logarithmic grid in P. The serving threshold was non-decreasing in m at all grid points except one adjacent pair, which we attribute to grid resolution; indexability was therefore checked numerically, not proved.

Results are total time-average error covariances obtained by propagating (4)–(5) under each schedule; for the linear-Gaussian model they carry no sampling error. The Monte-Carlo validation (eight seeds of 2×10^5 slots with sampled noise) and the lossy-link study (Bernoulli losses, eight seeds of 2.5×10^5 slots) are reported as means with 95% confidence intervals. Observation 1 was checked on the grid of Section 2.3.2; apparent violations in double precision were re-evaluated in 60-digit arithmetic and were at the resolution of the evaluation.

The measured case study uses the UCI Gas Sensor Array Drift Dataset [35,36]: sixteen metal-oxide sensors of four types, 13,910 measurements in ten batches over 36 months. Since no reference concentrations are released, each sensor's uncertainty is estimated by co-calibration: the sensor is predicted from the three other units of its type with a calibration fitted on batch 1 and then frozen, and the residual variance in later batches, normalised to the batch-1 baseline, is taken as its R. Only R is measured; the process dynamics are modelled.

3 RESULTS AND DISCUSSION

3.1 Comparison with baselines

In the configuration $R = (0.1,0.1,0.1,100)$, $a = 1$, MUD reduces total MSE by 11.5% relative to round robin, while its mean AoI rises from 1.50 to 1.96 slots. Greedy AoI and greedy Δ coincide with round robin in this configuration, and greedy covariance is worse than round robin (by 4.9–14.3% across the configurations tested) because it over-serves the noisiest sensor. That minimum age and minimum estimation error need not coincide is consistent with earlier findings [10,11,15]; the result here isolates the effect of instrument heterogeneity.

Because round robin and greedy rules are weak comparators, Table 3 includes the Whittle-index scheduler. With identical instruments every policy coincides with round robin, as Theorem 1(i) predicts. In the mixed configuration, where source 4 is stable with $R = 100$, the Whittle index also gives that source a share of 0.0000, i.e. a method derived without the information budget abandons the source identified by Proposition 3. In the two heterogeneous integrating configurations the Whittle index is better than MUD by 0.49% and 0.42%; MUD recovers 96% of the Whittle improvement over round robin. The remaining gap reflects that MUD fixes long-run rates while the Whittle index reacts to the current covariance. In exchange, the MUD allocation took tens of milliseconds, against about seven seconds to construct the index, and yields a static schedule.

Table 3: Total MSE against a Whittle-index scheduler, 6×10^4 slots. RR: round robin; Gr. P: greedy covariance.

Configuration	RR	Gr. P	Whittle	MUD
$R = (1,1,1,1)$, $a = 1$	9.3137	9.3137	9.3137	9.3137
$R = (0.1,0.1,0.1,100)$, $a = 1$	24.3926	27.8762	21.4726	21.5772
$R = (0.1,1,10,100)$, $a = 1$	29.6590	32.7617	27.1349	27.2492
$R = (0.1,0.1,0.1,100)$, $a_4 = 0.9$	9.7332	11.0701	8.5538	8.5538

3.2 The two branches across the parameter space

Figure 3 shows the optimal rate over the (a, R) plane at a fixed price. For $a \geq 1$ the rate increases with R (proved for $a^2 = 1$, observed numerically for $a > 1$); for $a < 1$ it peaks and then falls to zero in the hatched region. The dashed curve is (14). Bisecting on R until the rate collapses reproduces (14) to solver tolerance, from $R^\dagger = 5.93$ at $a = 0.85$ to 969.73 at $a = 0.97$; since both routes use the same model, this checks the derivation and implementation only.

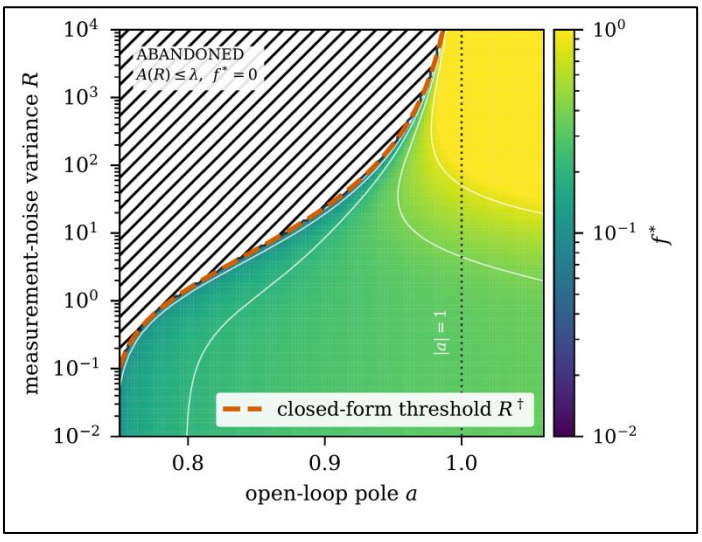

Figure 3: Optimal service rate over the (a, R) plane, $\lambda = 5$, $Q = 1$.

The size of the effect depends on the spread of instrument variances. In a sweep of the ratio between the best and worst instrument, the gain over round robin was under 1.4% for less than one decade, 10.8% at two decades and above 24% at three. Small spreads therefore yield small gains; the gains above apply to populations as heterogeneous as Table 1.

3.3 Probing floor and misreported variances

For a stable source ($a = 0.95$, $R = 1000$, threshold 209.5 at $\lambda = 4.91$) sharing the channel with three integrating sources ($a = 1, R = 1$), the unfloored optimum abandons it. A floor $\epsilon = 0.01$ caps its age at 100 slots at a cost of 0.25% in MSE (still 7.75% better than round robin); $\epsilon = 0.05$ costs 1.31%. When the instrument is recalibrated mid-run, the unfloored policy never re-admits it, whereas with $\epsilon = 0.01$ the improvement is detected after 50 slots on average.

Table 4 reports the effect of device 4 reporting $\hat{R}_4 = \rho R_4$ in the second configuration of Table 3, with excess MSE measured against the continuous-rate optimum (relative to which round robin is 13.93% worse). Moderate errors are cheap, but the response is asymmetric: under-reporting loses at most what uncertainty-blind scheduling loses, whereas a large over-report captures most of the channel and performs worse than round robin. A network manager using reported variances should therefore bound them, for example by the instrument's accuracy class or calibration certificate [32].

Table 4: Cost of a misreported variance, $N = 4$, $a = 1$.

$\rho = \hat{R}_4/R_4$	Excess MSE	Share of device 4
0.01	13.09%	25.6%
0.10	5.84%	32.0%
0.50	0.60%	41.8%
1.00 (truthful)	0.00%	46.7%
2.00	0.62%	51.5%
10.0	7.29%	62.5%
100	32.19%	75.7%

3.4 Accuracy of the relaxation and the model

The Monte-Carlo simulation gives 21.609 ± 0.144 for MUD against a Riccati prediction of 21.577, and 24.446 against 24.393 for round robin, both inside the confidence intervals. The realised improvement, 11.6%, is below the 12.2% of the continuous-rate optimum; the difference is due to integer slot grants. The online index (15) reaches shares within 4.9×10^{-4} of the offline rates after 3000 slots, and the tracking error decayed approximately as $O(1/k)$ over three decades; this rate is empirical. The gain depends on the per-device rate rather than on N: on one channel it falls from 11.3% at $N = 2$ to 0.95% at $N = 32$, but stays between 5.9% and 8.7% up to $N = 32$ when $B/N = 0.25$.

3.5 Lossy links

Against simulated Bernoulli losses, the effective-rate surrogate underestimates the cost by 0.8% at $p = 0$, 4.7% at $p = 0.1$ and 24.0% at $p = 0.5$. We therefore use it only to compute allocations and evaluate them by simulation. Table 5 separates link knowledge from instrument knowledge for one configuration: link-aware allocation alone recovers little, whereas uncertainty-aware allocation alone recovers nearly all of the combined gain. Since link quality is already monitored by network managers while instrument uncertainty usually is not, this is the dimension the proposed policy adds.

Table 5: Gain decomposition, $\boldsymbol{p} = (\mathbf{0.30}, \mathbf{0.10}, \mathbf{0.05}, \mathbf{0.20})$, $\boldsymbol{R} = (\mathbf{0.1}, \mathbf{0.1}, \mathbf{0.1}, \mathbf{100})$; simulated total MSE, mean $\pm$ 95% CI.

Scheduler	Simulated MSE	Gain vs fair
Freshness-fair	29.557 ± 0.015	–
Link-aware only	29.437 ± 0.017	0.40%
Uncertainty-aware only	26.838 ± 0.022	9.20%
Both (proposed)	$\mathbf{26.668 \pm 0.025}$	9.77%

3.6 Measured sensor drift

Table 6 summarises the co-calibration uncertainty after 36 months. Median R differs by about 60× between sensor types and grows by a median factor of 193 over the period. With these measured trajectories and modelled dynamics,

MUD reduces total MSE by 3.50% relative to freshness-fair allocation. This is smaller than in the synthetic case, as expected from Section 3.2: the spread is under two decades between types and as small as 1.5× within a type. When the monotone component of each trajectory is extracted by isotonic regression, the reduction is 3.18%, and the set of abandoned (floored) sensors grows from zero to ten as the allocation and price are recomputed batch by batch. Because isotonic fitting imposes monotonicity, this second analysis is semi-synthetic. Figure 4 shows the acetone channel together with a reference threshold $R^{\dagger} = 12.62$ computed at a fixed price; the operative threshold moves with the price.

Table 6: Measured uncertainty by sensor type at 36 months (median *R* normalised to the batch-1 baseline).

Sensor type	Median *R*	Within-type spread
TGS2600	105.1	14.4×
TGS2602	2062.8	3.9×
TGS2610	34.4	1.5×
TGS2620	259.9	2.0×

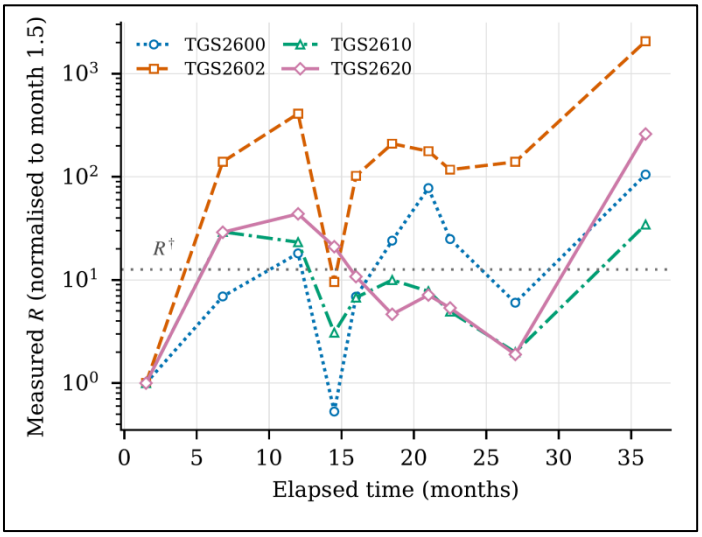

Figure 4: Measured uncertainty growth across ten batches, UCI gas sensor array, acetone channel; dashed line: reference threshold $R^{\dagger} = 12.62$.

Limitations

The analysis is scalar and applies to the continuous-rate relaxation (6); integer schedules differed from it by about 1–2% in the cases tested. Convexity, and hence Proposition 3 and the uniqueness of the optimum, is proved only for $a^2 = 1$ and otherwise rests on Observation 1. Theorem 1(ii) does not hold for stable sources. The lossy-link model discards rather than retransmits and its surrogate is optimistic; delay and correlated fading are not modelled. The Whittle-index scheduler retains a small advantage, so the rate-based policy is not optimal over all schedules, and stronger baselines, such as optimal periodic schedules from the sensor-scheduling literature [28], were not evaluated. The evaluation is synthetic except for the drift study, which uses measured uncertainty but modelled processes; its sixteen channels pair off with correlations above 0.994, co-calibration cannot detect drift common to a sensor type, and the batches are unevenly spaced. Validation on an instrumented TSCH testbed remains future work.

4 CONCLUSION

Slot allocation in a shared industrial wireless channel should not treat all measurements as equally informative. When instruments differ, the appropriate response depends on the monitored process: integrating sources warrant more bandwidth as their instruments degrade, whereas for self-regulating sources the value of a slot is capped by an instrument information budget, so sufficiently degraded instruments can be served only at a minimal probing rate. A rate-based allocation that uses instrument variance requires only a scalar search and a static slot pattern, and in the cases studied it came close to a state-dependent index scheduler. The practical significance is that instrument uncertainty, which is increasingly available from calibration records and online estimation, could be used by network managers alongside link quality. Remaining questions include a convexity proof for general poles, vector-state models, retransmission-based lossy links, incentive-compatible variance reporting and validation on an instrumented testbed.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Acknowledgments

The authors thank their institutions for their support.

Disclosure of conflict of interest

The authors declare that they have no known competing financial or non-financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

[1] International Electrotechnical Commission. Industrial networks – Wireless communication network and communication profiles – WirelessHART. IEC 62591:2016. Geneva: IEC; 2016.

[2] Thubert P. An architecture for IPv6 over the TSCH mode of IEEE 802.15.4 (6TiSCH). RFC 9030. Internet Engineering Task Force; 2021. doi:10.17487/RFC9030

[3] Kaul S, Yates R, Gruteser M. Real-time status: how often should one update? In: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM); 2012. p. 2731-5. doi:10.1109/INFCOM.2012.6195689

[4] Kaul S, Gruteser M, Rai V, Kenney J. Minimizing age of information in vehicular networks. In: Proceedings of the IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON); 2011. p. 350-8. doi:10.1109/SAHCN.2011.5984917

[5] Sun Y, Uysal-Biyikoglu E, Yates RD, Koksal CE, Shroff NB. Update or wait: how to keep your data fresh. IEEE Trans Inf Theory. 2017;63(11):7492-508. doi:10.1109/TIT.2017.2735804

[6] Kadota I, Sinha A, Uysal-Biyikoglu E, Singh R, Modiano E. Scheduling policies for minimizing age of information in broadcast wireless networks. IEEE/ACM Trans Netw. 2018;26(6):2637-50. doi:10.1109/TNET.2018.2873606

[7] Bedewy AM, Sun Y, Kompella S, Shroff NB. Optimal sampling and scheduling for timely status updates in multi-source networks. IEEE Trans Inf Theory. 2021;67(6):4019-34. doi:10.1109/TIT.2021.3060387

[8] Tripathi V, Modiano E. A Whittle index approach to minimizing functions of age of information. IEEE/ACM Trans Netw. 2024;32(6):5144-58. doi:10.1109/TNET.2024.3452006

[9] Yates RD, Sun Y, Brown DR, Kaul SK, Modiano E, Ulukus S. Age of information: an introduction and survey. IEEE J Sel Areas Commun. 2021;39(5):1183-210. doi:10.1109/JSAC.2021.3065072

[10] Sun Y, Polyanskiy Y, Uysal E. Sampling of the Wiener process for remote estimation over a channel with random delay. IEEE Trans Inf Theory. 2020;66(2):1118-35. doi:10.1109/TIT.2019.2937336

[11] Ornee TZ, Sun Y. Sampling and remote estimation for the Ornstein-Uhlenbeck process through queues: age of information and beyond. IEEE/ACM Trans Netw. 2021;29(5):1962-75. doi:10.1109/TNET.2021.3078137

[12] Pu C, Yang H, Wang P, Dong C. AoI-bounded scheduling for industrial wireless sensor networks. Electronics. 2023;12(6):1499. doi:10.3390/electronics12061499

[13] Liu CF, Bennis M. Taming the tail of maximal information age in wireless industrial networks. IEEE Commun Lett. 2019;23(12):2442-6. doi:10.1109/LCOMM.2019.2940965

[14] Chen K, Ding L, Yang F, Qian L. MSE minimized scheduling for multiple-source remote estimation with AoI constraints in IWSN. In: Proceedings of the International Conference on Wireless Communications and Signal Processing (WCSP); 2020. p. 1119-24. doi:10.1109/WCSP49889.2020.9299851

[15] Maatouk A, Kriouile S, Assaad M, Ephremides A. The age of incorrect information: a new performance metric for status updates. IEEE/ACM Trans Netw. 2020;28(5):2215-28. doi:10.1109/TNET.2020.3005549

[16] Zheng X, Zhou S, Niu Z. Urgency of information for context-aware timely status updates in remote control systems. IEEE Trans Wirel Commun. 2020;19(11):7237-50. doi:10.1109/TWC.2020.3009881

[17] Chen G, Liew SC, Shao Y. Uncertainty-of-information scheduling: a restless multi-armed bandit framework. IEEE Trans Inf Theory. 2022;68(9):6151-73. doi:10.1109/TIT.2022.3177891

[18] Soleymani T, Baras JS, Hirche S. Value of information in feedback control: quantification. IEEE Trans Automat Contr. 2022;67(7):3730-7. doi:10.1109/TAC.2021.3113472

[19] Soleymani T, Baras JS, Hirche S, Johansson KH. Value of information in feedback control: global optimality. IEEE Trans Automat Contr. 2023;68(6):3641-7. doi:10.1109/TAC.2022.3194125

[20] Ayan O, Vilgelm M, Klügel M, Hirche S, Kellerer W. Age-of-information vs. value-of-information scheduling for cellular networked control systems. In: Proceedings of the ACM/IEEE International Conference on Cyber-Physical Systems (ICCPS); 2019. p. 109-17. doi:10.1145/3302509.3311050

[21] Kountouris M, Pappas N. Semantics-empowered communication for networked intelligent systems. IEEE Commun Mag. 2021;59(6):96-102. doi:10.1109/MCOM.001.2000604

[22] Uysal E, Kaya O, Ephremides A, Gross J, Codreanu M, Popovski P, et al. Semantic communications in networked systems: a data significance perspective. IEEE Netw. 2022;36(4):233-40. doi:10.1109/MNET.106.2100636

[23] Joint Committee for Guides in Metrology. Evaluation of measurement data – Guide to the expression of uncertainty in measurement. JCGM 100:2008. Sèvres: BIPM; 2008. Available from: https://www.bipm.org/en/committees/jc/jcgm/publications

[24] Sinopoli B, Schenato L, Franceschetti M, Poolla K, Jordan MI, Sastry SS. Kalman filtering with intermittent observations. IEEE Trans Automat Contr. 2004;49(9):1453-64. doi:10.1109/TAC.2004.834121

[25] Lipsa GM, Martins NC. Remote state estimation with communication costs for first-order LTI systems. IEEE Trans Automat Contr. 2011;56(9):2013-25. doi:10.1109/TAC.2011.2139370

[26] Wu J, Jia QS, Johansson KH, Shi L. Event-based sensor data scheduling: trade-off between communication rate and estimation quality. IEEE Trans Automat Contr. 2013;58(4):1041-6. doi:10.1109/TAC.2012.2215253

[27] Han D, Mo Y, Wu J, Weerakkody S, Sinopoli B, Shi L. Stochastic event-triggered sensor schedule for remote state estimation. IEEE Trans Automat Contr. 2015;60(10):2661-75. doi:10.1109/TAC.2015.2406975

[28] Wu S, Ren X, Dey S, Shi L. Optimal scheduling of multiple sensors over shared channels with packet transmission constraint. Automatica. 2018;96:22-31. doi:10.1016/j.automatica.2018.06.019

[29] Wang J, Ren X, Mo Y, Shi L. Whittle index policy for dynamic multichannel allocation in remote state estimation. IEEE Trans Automat Contr. 2020;65(2):591-603. doi:10.1109/TAC.2019.2912492

[30] Gupta V, Chung TH, Hassibi B, Murray RM. On a stochastic sensor selection algorithm with applications in sensor scheduling and sensor coverage. Automatica. 2006;42(2):251-60. doi:10.1016/j.automatica.2005.09.016

[31] Harris P, Østergaard PF, Tabandeh S, Söderblom H, Kok G, van Dijk M, et al. Measurement uncertainty evaluation for sensor network metrology. Metrology. 2025;5(1):3. doi:10.3390/metrology5010003

[32] Mustapää T, Nikander P, Hutzschenreuter D, Viitala R. Metrological challenges in collaborative sensing: applicability of digital calibration certificates. Sensors. 2020;20(17):4730. doi:10.3390/s20174730

[33] Martins A, Fonseca I, Farinha JT, et al. Online monitoring of sensor calibration status to support condition-based maintenance. Sensors. 2023;23(5):2402. doi:10.3390/s23052402

[34] Bisdikian C, Kaplan LM, Srivastava MB. On the quality and value of information in sensor networks. ACM Trans Sens Netw. 2013;9(4):48. doi:10.1145/2489253.2489265

[35] Vergara A, Vembu S, Ayhan T, Ryan MA, Homer ML, Huerta R. Chemical gas sensor drift compensation using classifier ensembles. Sens Actuators B Chem. 2012;166-167:320-9. doi:10.1016/j.snb.2012.01.074

[36] Fonollosa J, Rodriguez-Lujan I, Huerta R. Chemical gas sensor array dataset. Data Brief. 2015;3:85-9. doi:10.1016/j.dib.2015.01.003