Age of Anomalies for Energy-Constrained Networked Systems

Houman Asgari
Technical University of Munich
Munich, Germany
houman.asgari@tum.de

Leonardo Badia
University of Padova
Padova, Italy
leonardo.badia@unipd.it

Andrea Munari
German Aerospace Center (DLR)
Weßling, Germany
andrea.munari@dlr.de

***Abstract*—This paper investigates a finite-horizon transmission control problem for a warning-based cyber-physical system with three operating modes: stable, vulnerable, and critical. In this setting, a sensor monitors the system and communicates updates to a controller under a finite-capacity battery with energy harvesting. At each time slot, the sensor decides whether to remain idle, transmit a compressed update, or transmit a dense update, with different energy costs and different stabilization effectiveness. Our objective is to design a finite-horizon transmission control policy under the battery dynamics. Performance is measured by a state-aware metric that captures both the onset of the critical mode and the duration of remaining in it. We formulate the problem as a finite-horizon Markov decision process and solve it via dynamic programming. Numerical results illustrate a time-varying threshold-type structure in the optimal policy, quantify when simple mode-based heuristics are near-optimal, and highlight the roles of critical-mode persistence and battery capacity.**

***Index Terms*—Cyber-physical systems, energy harvesting, battery-constrained communication, Markov decision processes, finite-horizon optimal control**

I. Introduction

Timely information exchange is fundamental in networked control and cyber-physical systems, where remote sensors communicate status updates to a controller. Communication delays, packet losses, and energy limitations can degrade control performance and lead to unsafe or unstable operation [1]. This challenge is especially pronounced in energy-constrained sensing platforms, where communication decisions must balance stabilization effectiveness against limited energy availability.

A large body of recent work has focused on the concept of *Age of Information (AoI)* as a metric for quantifying information freshness in networked systems [2]. AoI has been studied extensively in the context of scheduling, sampling, and resource allocation under communication constraints [3].

This latter aspect of considering a limited number of transmission opportunities has been connected to energy constraints in [4], and developed by subsequent papers to consider multiple sensors [5] or multiple sources [6]. More recent contributions also investigate optimal online algorithms [7] or the inclusion of additional elements such as eavesdropping [8].

However, in many cyber-physical applications, the primary concern is not merely how fresh the information is, but rather how long the system has remained in an undesirable or unsafe operating condition. This aligns with a broader rethinking of networked system design around the significance of information relative to the purpose of the data exchange, moving beyond classical rate- and delay-centric paradigms [9]. This observation has motivated the development of state-dependent age metrics, such as the *Age of Incorrect Information (AoII)*, which penalizes outdated or incorrect information more heavily when the system state is critical [10]. Complementary to this line of work, scheduling strategies that prioritize updates based on the operational importance of the underlying source state have been studied for discrete-state Markov models in [11]. Recently, the *Age of Anomalies (AoA)* has been proposed as a control-oriented performance metric that explicitly captures both the onset and the persistence of instability in networked control systems [12].

Another important characteristic of many real-world systems is that instability rarely arises abruptly. Instead, degradation toward unsafe operation is often preceded by detectable warning signs or intermediate conditions. This observation motivates system models that distinguish between *stable*, *vulnerable*, and *unstable* operating modes. In parallel, energy harvesting and energy-constrained communication have received significant attention in remote monitoring and control systems [13], [14]. In many practical scenarios, energy availability is tightly coupled to system operation: energy may be harvested only under favorable conditions, while degraded or critical states limit or prevent recharging. This coupling introduces a nontrivial tradeoff between early intervention to prevent instability and conserving energy for potentially more severe future events [15].

A. Munari would like to acknowledge the support of the German Federal Ministry of Research, Technology, and Space (BMFTR) with the xG-RIC project as part of the research program Communication Systems "Souverän. Digital. Vernetzt." (grant number 16KIS2429K).

Understanding how this tradeoff should be managed is essential for designing robust and implementable control policies. A representative example is vibration-based condition monitoring of industrial machinery. In such systems, machine vibrations can be used to power wireless sensors through piezoelectric or electromagnetic energy harvesting [16]. The harvested energy depends on the operating regime and vibration profile, while the urgency of communication increases as the monitored process moves toward unsafe operation.

From a control perspective, the problem naturally leads to a Markov decision process (MDP), where transmission actions affect both the system dynamics and battery state [3], [5]. Although dynamic programming yields optimal finite-horizon policies, these policies are time-varying and depend on multiple state variables, which complicates implementation on resource-constrained devices. This raises the question of when simple heuristics can achieve near-optimal performance.

This paper addresses this question in the context of AoA-aware control with warning-based system dynamics and energy harvesting constraints. Building on the binary mode, energy-unconstrained, infinite-horizon setting of [12], we extend the framework along three key dimensions. First, we introduce a three-mode system with *stable*, *vulnerable*, and *unstable* operating modes. The vulnerable mode captures warning dynamics and creates a nontrivial tradeoff between early intervention and energy conservation. Second, we incorporate a finite-capacity battery with state-dependent harvesting [6]. As a result, the optimal decision boundary becomes two-dimensional in the (AoA, battery) space. Third, we adopt a finite-horizon formulation, which models operation over supervision windows between external interventions or system resets [17]. In this model, the vulnerable mode represents an intermediate regime in which the system is not yet unstable and may still recover, but the energy harvesting is temporarily unavailable. This reflects practical scenarios in which the energy source is tied to normal system operation: for instance, a vibration harvester mounted on machinery draws energy from regular operational vibrations, which are absent or irregular during warning or fault conditions [18]. AoA accumulates only during operation in the unstable mode and resets otherwise. Transmission actions differ in both energy cost and stabilization effectiveness.

The main contributions of this paper are as follows:

- We characterize a finite-horizon AoA-aware optimal controller, revealing a time-varying double-threshold behavior in AoA and battery level shaped by instability persistence, stochastic harvesting, and time-to-go.
- We identify instability persistence as the dominant parameter governing when simple policies suffice: when instability is transient, simple mode-based heuristics achieve near-optimal performance; when instability is persistent, energy-aware control becomes critical.
- We demonstrate that battery provisioning can yield larger reductions in average AoA than switching from the best heuristic to the fully optimal policy, though with diminishing returns beyond moderate battery sizes.

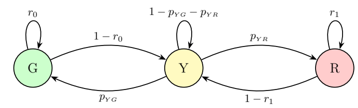

Fig. 1. Three-mode source model.

The remainder of the paper is organized as follows. Section II introduces the system model and problem formulation. Section III analyzes the structure of the finite-horizon optimal policy. Section IV presents numerical experiments evaluating the impact of horizon length, instability persistence, and battery capacity, and compares optimal and heuristic policies. Section V concludes the paper and outlines directions for future work.

II. System Model and Problem Formulation

We model the system as a discrete-time controlled stochastic process evolving over time slots $t \in \mathbb{Z}_{\geq 0}$. At each time t, the system operates in one of three modes, denoted by $\sigma_t \in \mathcal{S} \triangleq \{\mathrm{G}, \mathrm{Y}, \mathrm{R}\}$, where G, Y, and R represent the stable, vulnerable, and unstable modes, respectively.

The system state at time t is defined as

$$x_t \triangleq (s_t, \sigma_t, b_t) \tag{1}$$

where $s_t \in \mathbb{Z}_{\geq 0}$ denotes the AoA, which will be formally introduced shortly, σ_t is the system mode, and $b_t \in \mathcal{B} \triangleq \{0, 1, \ldots, B\}$ represents the available battery energy. At each time t, the controller selects a transmission action

$$a_t \in \mathcal{A}(b_t) \tag{2}$$

where the action set depends on the available battery energy and is given by

$$\mathcal{A}(b) \triangleq \{a \in \{\mathrm{I}, \mathrm{C}, \mathrm{D}\} : e(a) \leq b\}. \tag{3}$$

The actions correspond to *Idle* (I), *Compressed* (C), and *Dense* (D) transmission, with energy costs

$$e(\mathrm{I}) = 0, \qquad e(\mathrm{C}) = e_{\mathrm{C}}, \qquad e(\mathrm{D}) = e_{\mathrm{D}}, \tag{4}$$

where $0 < e_{\mathrm{C}} < e_{\mathrm{D}}$.

A. System Mode Dynamics

The system mode σ_t evolves as a Markov chain over the three macroscopic modes $\mathcal{S}$, as illustrated in Fig. 1. For each action $a \in \{\mathrm{I}, \mathrm{C}, \mathrm{D}\}$, let $P^{(a)}(\sigma' \mid \sigma)$ denote the transition probability from mode σ to mode σ'. Under the idle action ($a = \mathrm{I}$), the system follows its natural (uncontrolled) dynamics specified by the transition structure in Fig. 1, with transition matrix

$$P^{\mathrm{nat}} = \begin{bmatrix} r_0 & 1-r_0 & 0 \\ p_{YG} & 1-p_{YG}-p_{YR} & p_{YR} \\ 0 & 1-r_1 & r_1 \end{bmatrix} \tag{5}$$

where entry (i,j) gives $P^{\mathrm{nat}}(\sigma_j \mid \sigma_i)$ and the rows correspond to G, Y, R, respectively. Here, $r_0 \in (0,1)$ denotes the self-transition probability of the G mode, $r_1 \in (0,1)$ denotes the self-transition (persistence) probability of the R mode, $p_{YG} \in (0,1)$ is the Y-to-G recovery probability, and $p_{YR} \in (0,1)$ is the Y-to-R deterioration probability. To account for uplink channel erasures, a transmitted update is successfully received with probability $\rho \in (0,1)$. Otherwise, the transmission has no stabilizing effect and the system follows the natural dynamics. For transmission actions $a \in \{\mathrm{C}, \mathrm{D}\}$, conditioned on successful reception, the system may undergo action-induced stabilization along the admissible transitions of the three-mode model. Compressed transmission improves the mode by one level (R $\to$ Y or Y $\to$ G) with probability $p_{\mathrm{C}} \in (0,1)$, while dense transmission enables the chance of a direct transition to the stable mode G (R $\to$ G or Y $\to$ G) with probability $p_{\mathrm{D}} \in (0,1)$. All remaining probability mass is assigned to the natural transitions captured by P^{nat}.

B. Battery Dynamics

Energy harvesting occurs only when the system is in the G mode. Let $h_t \in \{0,1\}$ denote the harvested energy at time t, where

$$\mathbb{P}(h_t = 1 \mid \sigma_t) = \begin{cases} p_{\mathrm{ch}}, & \sigma_t = \mathrm{G}, \\ 0, & \sigma_t \in \{\mathrm{Y}, \mathrm{R}\}. \end{cases} \tag{6}$$

The battery evolves according to

$$b_{t+1} = \min\{B,\ b_t - e(a_t) + h_t\}. \tag{7}$$

C. AoA Dynamics

The AoA evolves deterministically as

$$s_{t+1} = \begin{cases} 0, & \sigma_{t+1} \in \{\mathrm{G}, \mathrm{Y}\}, \\ s_t + 1, & \sigma_{t+1} = \mathrm{R}. \end{cases} \tag{8}$$

The above dynamics define a controlled Markov process on the state space

$$\mathcal{X} = \mathbb{Z}_{\geq 0} \times \mathcal{S} \times \mathcal{B}. \tag{9}$$

Fig. 2 illustrates a sample path of the AoA alongside the corresponding mode sequence. The AoA increases linearly during consecutive R episodes and resets to zero upon any transition to G or Y.

D. Mean AoA Criterion

To minimize the AoA over time, we formulate the transmission control problem as an MDP, where the per-stage cost is defined as the instantaneous AoA:

$$c(x_t, a_t) \triangleq s_t. \tag{10}$$

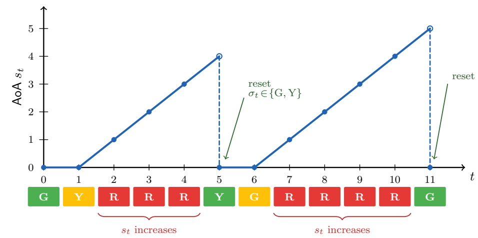

Fig. 2. Sample evolution of the AoA. The colored ribbon shows the system mode σ_t, while the sawtooth curve traces s_t: it ramps up during R episodes and resets to zero when the mode returns to G or Y.

Our primary objective is the finite-horizon formulation, which captures operation over a supervision window of length T. The optimal expected cumulative AoA is

$$V_T^{\mathrm{mean}}(x) = \inf_{\pi} \mathbb{E}_x^{\pi}\left[\sum_{t=0}^{T-1} c(x_t, a_t)\right], \tag{11}$$

where π ranges over admissible policies. For completeness, we also state the corresponding infinite-horizon average-cost objective, which serves as a limiting benchmark as $T \to \infty$:

$$g_{\mathrm{mean}}^{\star} = \inf_{\pi} \limsup_{T\to\infty} \frac{1}{T} \mathbb{E}_x^{\pi}\left[\sum_{t=0}^{T-1} c(x_t, a_t)\right]. \tag{12}$$

Beyond its role as a per-stage cost, AoA admits a backlog interpretation. It increases while the system operates in the critical mode R and resets when the system recovers to G or Y. Thus, AoA measures the accumulated duration of unresolved instability, with larger values indicating prolonged failure to restore stable operation.

This interpretation is analogous to classical Lyapunov-drift arguments in stochastic stability theory, where the boundedness of a backlog-like state variable indicates system stability. In this sense, maintaining bounded AoA corresponds to preventing the accumulation of instability over time [19].

III. Finite-Horizon Optimal Control

We first consider the finite-horizon formulation over a supervision window of length T [17]. The corresponding optimal policy serves primarily as a performance benchmark against which simpler heuristic strategies are evaluated in Section IV. Although the policy can be computed exactly via dynamic programming, its dependence on full state information and remaining time-to-go makes direct implementation challenging in resource-constrained settings.

The objective is to minimize the expected cumulative AoA over the horizon, subject to the system dynamics and energy constraints introduced in Section II. Let $V_t(x)$ denote the optimal expected remaining cost from time t to the end of the horizon when the system is in state x. The terminal condition is $V_T(x) = 0$ for all $x \in \mathcal{X}$.

For $t = T-1, T-2, \ldots, 0$, the value function satisfies the Bellman recursion

$$V_t(x) = \min_{a \in \mathcal{A}(b)} \Bigg\{ c(x_t, a) + \sum_{x' \in \mathcal{X}} P(x' \mid x_t, a)\, V_{t+1}(x') \Bigg\}. \tag{13}$$

The optimal finite-horizon policy is time-dependent and is obtained by selecting, at each time t, an action attaining the minimum in the above recursion. Since both the state and action spaces are finite, the optimal value function and policy can be computed exactly via backward induction.

IV. Discussion and Simplified Policies

A. Finite-Horizon Policy Structure

We begin by illustrating the structure of the optimal finite-horizon policy through three representative snapshots of the action-state map. Unless stated otherwise, the system parameters used throughout the simulations are as follows. The decision horizon is set to $T = 150$, and the AoA is truncated at 50 for computational tractability. The battery capacity is $B = 7$, with initial battery level $b_0 = 3$. Compressed and dense transmissions consume $e_{\mathrm{C}} = 1$ and $e_{\mathrm{D}} = 3$ units of energy, respectively. A transmitted update is successfully received with probability $\rho = 0.9$, while the success probabilities associated with compressed and dense transmissions are $p_{\mathrm{C}} = 0.3$ and $p_{\mathrm{D}} = 0.7$.

The natural evolution of the system mode follows the three-state Markov chain described in Section II, with transition probabilities $r_0 = 0.9$, $r_1 = 0.85$, $p_{\mathrm{YG}} = 0.1$, and $p_{\mathrm{YR}} = 0.2$. Energy harvesting is available only in the stable mode, with harvesting probability $p_{\mathrm{ch}} = 0.5$.

Fig. 3 depicts the optimal action-state maps for the unstable (R) mode at two representative decision times $t \in \{20, 130\}$, corresponding to an early stage and a time close to the end of the horizon. In each plot, the horizontal axis represents the battery level and the vertical axis represents the instantaneous AoA, while the color of each cell indicates the selected transmission action.

Several qualitative features emerge from these snapshots. First, the optimal policy exhibits a pronounced threshold-like behavior in the AoA dimension: for a fixed battery level, the controller remains idle for small AoA values and switches to active transmission once the AoA exceeds a state-dependent threshold. Second, the location of this threshold depends on the available battery energy, with more aggressive actions selected at higher battery levels. Finally, the thresholds evolve over time: near the end of the horizon the policy becomes less conservative and favors transmission at smaller AoA

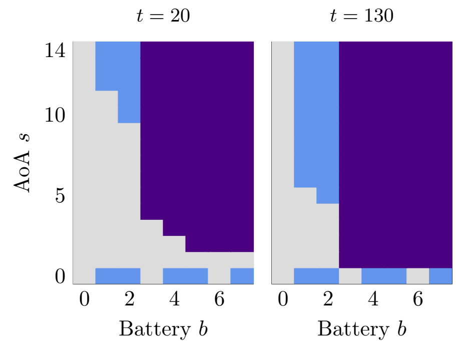

Fig. 3. Optimal policy at $t = 20$ (left) and $t = 130$ (right) in R mode. Gray: IDLE, blue: COMPRESSED, dark purple: DENSE. As the horizon approaches, thresholds decrease and the policy becomes more aggressive.

values, whereas earlier in the horizon the controller tolerates longer instability episodes before intervening.

While the overall structure is clearly threshold-like, the decision boundaries are not perfectly monotone or staircase-shaped [5], [20]. Small local irregularities appear near the transition regions between actions. This behavior is inherent to the finite-horizon formulation, as the optimal action depends jointly on the AoA, the battery level, and the remaining time-to-go. The interaction between energy availability, stochastic harvesting, and discrete state dynamics naturally leads to such fine-grained variations. These effects do not contradict the threshold interpretation; rather, they reflect the delicate structure of the optimal finite-horizon solution.

B. Effect of Horizon Length and Asymptotic Behavior

We next investigate the impact of the decision horizon on the optimal finite-horizon performance. Using the same system parameters as in the previous experiment, we solve the finite-horizon problem for horizons ranging from $T = 20$ to $T = 1000$. For each horizon, we report the optimal average AoA, defined as the expected total AoA normalized by the horizon length.

Fig. 4 shows the optimal average AoA as a function of the decision horizon. The figure also includes the infinite-horizon average-cost optimum $g^\star \approx 1.154$, computed independently via relative value iteration.

Two observations are worth highlighting. First, for short horizons, the finite-horizon average AoA is clearly below $g^\star$. At $T = 30$, the value is much smaller than the infinite-horizon benchmark, and a visible gap remains even at $T = 150$. This reflects the limited operating window: when only a small number of slots remain, the controller can afford to tolerate instability for longer, since its future cost impact is limited. This is the regime where the finite-horizon formulation is most relevant.

Second, the gap decreases as the horizon grows. By $T \approx 200$, the finite-horizon average AoA is already close to $g^\star$, and the remaining difference shrinks gradually for

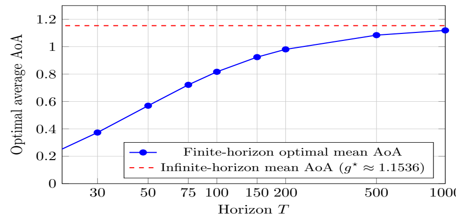

Fig. 4. Optimal average AoA as a function of the decision horizon T.

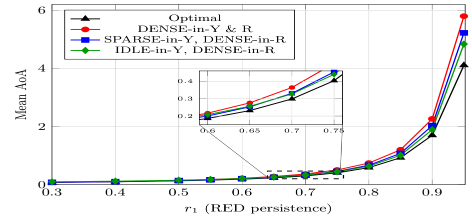

Fig. 5. AoA for different policies as a function of the Red persistence probability r_1 ($T = 150$).

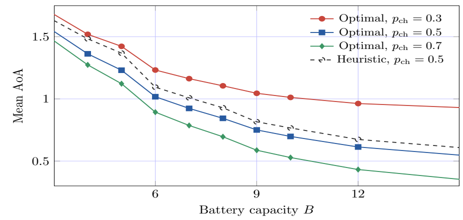

Fig. 6. Optimal average AoA vs. harvest probability for varying battery capacity ($r_1 = 0.85$).

larger T. This observation is consistent with average-cost optimal control theory [21], which predicts that

$$V_T(x) = Tg^\star + h(x) + o(1),$$

where $g^\star$ denotes the optimal long-run average AoA and $h(x)$ is the bias function encoding the state-dependent transient cost.

C. When Do Simple Policies Suffice?

The finite-horizon action-state maps in Fig. 3 show that, although the optimal policy has a clear threshold-like structure, its dependence on the AoA, battery level, and remaining time-to-go is nontrivial. In particular, the decision boundaries vary over time, resulting in a policy map that can be costly to store and execute on resource-constrained devices [20].

These considerations motivate the following question: *to what extent can simple, rule-of-thumb policies achieve near-optimal performance, and under which system conditions are such simplifications justified?* To address this question, we compare the optimal finite-horizon policy with several heuristic strategies that rely on fixed action rules based solely on the system mode, thereby eliminating explicit dependence on the AoA or the remaining horizon. Specifically, we consider the following three mode-based heuristics:

- *DENSE-in-Y & R:* Transmit densely whenever the system is in Y or R, and remain idle in G. This is the most aggressive strategy. It reflects the intuition that early intervention should reduce the time spent in the unstable mode.
- *SPARSE-in-Y, DENSE-in-R:* Transmit with compressed updates in Y and dense updates in R, while remaining idle in G. This strategy balances early reaction and energy conservation. It uses the cheaper action in the warning phase and reserves the stronger action for the critical mode.
- *IDLE-in-Y, DENSE-in-R:* Remain idle in Y and transmit densely only in R, while remaining idle in G. This is the most conservative strategy. It preserves energy during the vulnerable phase and uses it only once the system reaches the unstable mode.

Fig. 5 indicates that a clear regime distinction emerges as r_1, i.e., instability persistence, varies. For small values of r_1, all considered heuristic policies exhibit performance very close to the optimal solution. In this regime, episodes of unstable behavior are short-lived and are naturally resolved with high probability. As a result, the precise allocation of energy across Y and R modes has limited impact, and simple mode-based policies suffice.

As r_1 increases, however, the heuristics begin to separate. When R becomes more persistent, instability episodes last longer, and the consequences of suboptimal energy allocation become more pronounced. In this regime, the policy that remains idle in the Y mode and transmits densely in the R mode consistently outperforms the other simple strategies. This reflects the increased value of conserving energy during the vulnerable phase in order to sustain intervention throughout prolonged periods in the R mode.

These observations suggest that increased instability persistence shifts the dominant control challenge from reacting to transient events to managing sustained instability. Accordingly, policies that prioritize energy preservation in Y and concentrate intervention in R become increasingly advantageous as r_1 grows.

D. Battery Capacity Sensitivity

We next examine how battery capacity affects performance and whether increasing energy storage can compensate for policy simplification [4]. Since the previous experiment identified the *IDLE-in-Y, DENSE-in-R* heuristic as the best-performing simple policy in the high-persistence regime, we use it as the representative heuristic benchmark throughout this subsection. Figure 6 depicts the optimal finite-horizon average AoA as a function of the battery capacity B for several energy harvesting probabilities p_{ch}, together with the performance of the *IDLE-in-Y, DENSE-in-R* heuristic when $p_{\mathrm{ch}} = 0.5$.

As expected, for all harvesting probabilities, the optimal average AoA decreases monotonically as the battery capacity increases. However, the performance improvement exhibits diminishing returns beyond moderate battery sizes. For small B, the system is frequently energy-constrained, preventing timely intervention during unstable periods. As B increases, these energy shortages become less frequent, enabling more effective stabilization. Once the battery capacity is sufficiently large to buffer several unfavorable harvesting realizations, further increases provide only marginal benefit. In this regime, the dominant performance bottleneck shifts from energy storage to the energy harvesting process itself and the source's intrinsic dynamics.

An additional insight from Fig. 6 is that the performance improvement achieved by increasing battery capacity can exceed the improvement obtained by switching from a simple heuristic to the optimal policy. In particular, for both low and high battery regimes, enlarging the energy buffer yields a more pronounced reduction in average AoA than refining the control strategy while keeping the battery size fixed. This suggests that, from a system design perspective, investing in larger energy storage may provide greater practical benefit than implementing a fully state-aware optimal controller, especially when the heuristic already captures the dominant stabilization behavior.

Taken together, the numerical results highlight two dominant system parameters: the persistence of instability and the available energy budget. When instability is transient, simple policies suffice. When instability is persistent, energy allocation becomes critical. Across regimes, battery provisioning exerts an influence on performance, whereas policy refinement plays a secondary role except in intermediate energy-constrained settings.

V. Conclusion

We considered an energy-harvesting sensor that monitors a warning-based dynamical system and, at each time step, decides whether to stay idle or transmit (with different transmission options) under a finite battery constraint, with the goal of keeping the AoA small.

Through extensive simulations we depicted that the benefit of state-aware optimization is largely governed by the persistence of the instability: when instability episodes are short, simple mode-based heuristics perform close to optimal, whereas persistent instabilities make energy-aware allocation across vulnerable and unstable states critical. Moreover, increasing battery capacity can provide larger AoA reductions than replacing a reasonable heuristic with the fully optimal policy, except in the intermediate, energy-limited regime where refined decisions yield the greatest gains.

Future work can address model uncertainty, where the Markov model parameters governing the warning dynamics are not known a priori. This naturally leads to joint learning and control problem.

References

[1] K.-D. Kim and P. R. Kumar, "Cyber–physical systems: A perspective at the centennial," *Proc. IEEE*, vol. 100, no. Special Centennial Issue, pp. 1287–1308, 2012.

[2] S. Kaul, R. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *Proc. IEEE INFOCOM*, 2012, pp. 2731–2735.

[3] E. T. Ceran, D. Gündüz, and A. György, "Average age of information with hybrid ARQ under a resource constraint," *IEEE Trans. Wireless Commun.*, vol. 18, no. 3, pp. 1900–1913, 2019.

[4] X. Wu, J. Yang, and J. Wu, "Optimal status update for age of information minimization with an energy harvesting source," *IEEE Trans. Green Commun. Netw.*, vol. 2, no. 1, pp. 193–204, 2017.

[5] M. Mei, D. Ye, and J. Wei, "Multisensor power allocation for remote estimation over semi-Markov fading channels," *IEEE Trans. Control Netw. Syst.*, vol. 11, no. 3, pp. 1300–1309, 2023.

[6] E. Gindullina, L. Badia, and D. Gündüz, "Age-of-information with information source diversity in an energy harvesting system," *IEEE Trans. Green Commun. Netw.*, vol. 5, no. 3, pp. 1529–1540, 2021.

[7] Q. Lin, J. Su, and M. Chen, "Optimal algorithms for online age-of-information optimization in energy harvesting systems," *IEEE/ACM Trans. Netw.*, vol. 33, no. 6, pp. 3146–3161, 2025.

[8] F. Yuan, S. Tang, and D. Liu, "AoI-based transmission scheduling for cyber physical systems over fading channel against eavesdropping," *IEEE Internet Things J.*, vol. 11, no. 3, pp. 5455–5466, 2023.

[9] E. Uysal, O. Kaya, A. Ephremides, J. Gross, M. Codreanu, P. Popovski, M. Assaad, G. Liva, A. Munari, B. Soret, T. Soleymani, and K. H. Johansson, "Semantic communications in networked systems: A data significance perspective," *IEEE Network*, vol. 36, no. 4, pp. 233–240, 2022.

[10] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Trans. Netw.*, vol. 28, no. 5, pp. 2215–2228, 2020.

[11] J. Luo and N. Pappas, "Exploiting data significance in remote estimation of discrete-state markov sources," *IEEE Trans. Commun.*, vol. 74, pp. 4569–4582, 2026.

[12] S. Kriouile, M. Assaad, and T. Soleymani, "Semantics of anomalies in networked control," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2025, pp. 1–6.

[13] S. Ulukus, A. Yener, E. Erkip, O. Simeone, M. Zorzi, P. Grover, and K. Huang, "Energy harvesting wireless communications: A review of recent advances," *IEEE J. Sel. Areas Commun.*, vol. 33, no. 3, pp. 360–381, 2015.

[14] G. Stamatakis, N. Pappas, and A. Traganitis, "Control of status updates for energy harvesting devices that monitor processes with alarms," in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2019, pp. 1–6.

[15] M. Pagin, L. Badia, and M. Zorzi, "A markov game of age of information from strategic sources with full online information," in *Proc. IEEE Int. Conf. Commun.*, 2023, pp. 76–81.

[16] X. Tang, X. Wang, R. Cattley, F. Gu, and A. D. Ball, "Energy harvesting technologies for achieving self-powered wireless sensor networks in machine condition monitoring: A review," *Sensors*, vol. 18, no. 12, 2018.

[17] A. Munari and L. Badia, "What's my age of information again? the role of feedback in AoI optimization under limited transmission opportunities," *IEEE Trans. Commun.*, vol. 73, no. 11, pp. 10 495–10 508, 2025.

[18] F. Huet, V. Boitier, and L. Seguier, "Tunable piezoelectric vibration energy harvester with supercapacitors for WSN in an industrial environment," *IEEE Sensors J.*, vol. 22, no. 15, pp. 15 373–15 384, 2022.

[19] L. Wang, Q. Wang, H. H. Chen, and S. Zhou, "Age of information-oriented probabilistic link scheduling for device-to-device networks," in *Proc. Int. Symp. Modeling Optim. Mobile Ad Hoc Wirel. Netw. (WiOpt)*, 2024, pp. 201–208.

[20] C. Cavalagli, L. Badia, and A. Munari, "Reinforcement learning for age of information aware transmission policies in slotted ALOHA channels," in *Proc. Int. Symp. Wireless Commun. Syst. (ISWCS)*, 2024, pp. 1–6.

[21] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.