What About Peak Age? Average vs. Peak AoI Minimization in Finite-Horizon Scheduling

Beyza Türk
Dept. of Computer Engineering
Istanbul Technical University
Istanbul, Türkiye
turkbe22@itu.edu.tr

Leonardo Badia
Dept. of Information Engineering (DEI)
University of Padova
Padova, Italy
leonardo.badia@unipd.it

***Abstract*—Current communication standards typically emphasize latency, reliability, and throughput as the main performance metrics. However, a promising line of research is adopting age of information (AoI) as a more direct measure of data freshness, which is key for real-time applications and ambient sensing. In particular, industrial or mission-critical scenarios would likely require an AoI minimization over a finite horizon, e.g., corresponding to the duration of an intended observation window, under the constraint of a limited number of updates being exchanged. From the standpoint of a standard to be implemented in future communications and networking platforms, the choice would further be conflicted between minimizing the average or peak value of AoI. Anticipating this conundrum, we explore an optimization approach based on dynamic programming recursion, where we consider both average and peak AoI as possible objectives; we also evaluate one metric when the other is optimized. We further consider both independent and correlated errors. In all these cases, we are able to show that minimizing either AoI-related metric over a relatively short horizon often converges to similar threshold-based criteria. From a practical standpoint, this supports the insertion of either AoI-related metric in future standards, making further debates over average versus peak AoI minimization amount to just semantic distinctions with negligible practical impact.**

***Index Terms*—Industrial Internet of things; Age of Information; Optimization; IoT Design Options; Dynamic Programming.**

I. Introduction

For real-time communication exchange aimed at reporting status updates, the age of information (AoI) is defined as the time elapsed between the current time and the instant of generation of the last update successfully received [1]. Designed to capture how old the last received information about the state of a system can be, AoI is recognized by the research community as a useful metric for real-time communication and control in cyber-physical systems [2] and digital twins [3], i.e., whenever a virtual representation of the physical world is required to make an accurate and timely decision of acquisition and actuation.

With the possible advent of goal-oriented communications, decision and control mechanisms based on communication semantics can also become widespread, which also requires adherence between the real system and its digital representation [4]. Similarly, the expected diffusion of logical inference and reasoning through some large language models based on ambient observations makes it even more relevant that these virtual agents make their decisions based on information that is as up-to-date as possible [5], [6].

However, current communication standards do not use AoI as an explicit design criterion, such as IEEE 802.11, 802.15.4, 3GPP 5G NR [7], LoRaWAN [8], IETF CoAP/QUIC [9], or Time-Sensitive Networking (TSN) [10] all focus on similar yet different metrics such as latency or jitter. Nevertheless, there is growing recognition in the research community that real-time communications for status update reporting should be aligned with metrics related to freshness of information rather than traditional latency or throughput, and this can be expected to be reflected in future standardization efforts [11].

Some precursor contributions have already been proposed especially for the context of vehicular networks, where AoI was originally proposed [12]. Beyond that, AoI is gaining attention in many research proposals and whitepapers that, while pertaining to a pre-standardization area, suggest a possible integration in many contexts. For example, Ultra-Reliable Low-Latency Communications (URLLC) defined by 3GPP indirectly aligns with the principles of AoI, which led some AoI-aware scheduling proposals already suggested for 6G networks [13].

At the same time, it should be noted that the general idea of fresh status updates can be implemented for a finite-horizon scheduling [14] under two different declensions: minimization of average AoI or peak AoI [15]. The average AoI is the most common way to assess the freshness of reported information, to indicate that the system receives frequent and regular updates and operates with up-to-date information [16]. In contrast, Peak AoI represents a worst-case scenario, and it may be a preferred metric when the presence of outdated information creates a risk (e.g. emergency response systems, critical communication networks) [17], [18].

In general, these two metrics may be unaligned, since a single extremely delayed update leads to a high peak AoI but may be compensated for on average. For this reason, in this paper we discuss optimal AoI-aware scheduling for finite-horizon systems, comparing the minimization of average

This work has been supported by the Italian PRIN project 2022PNRR "DIGIT4CIRCLE," project code P2022788KK.

and peak AoI and seeking whether this results in significant difference. In both cases, we prove how optimization can be performed through a standard dynamic programming approach that takes advantage of the finiteness of the scheduling horizon through backward induction. We also consider different channel models affected by independent or correlated errors [19].

The results of our analysis actually show that the two objectives result in similar policies, both with a threshold-based structure. In addition, their performance is also very similar, where a minimum peak AoI scheduler obtains an average AoI that, while clearly not minimal, is still close to being optimum, and this also holds true in reverse. Finally, these conclusions are valid for both independent and correlated errors, with the specific conclusion that error correlation increases the average AoI (but similarly for both schedulers), while the peak AoI is largely unaffected [20].

This main finding serves as a useful guideline for indicating that the insertion of AoI-awareness would be useful in communication protocols for time-sensitive industrial or mission-critical applications. However, distinguishing between the precise AoI metric is likely to be of limited practical relevance, as the primary objective remains information freshness regardless of the specific formulation, rendering such academic distinctions of little interest in practical contexts.

The remainder of this paper is organized as follows. In Section II we discuss related work. The proposed methodology is presented in Section III. Section IV includes numerical evaluations and analysis. Finally, Section V concludes the paper and outlines future work.

II. Related Work

The introduction of AoI allows for a quantitative evaluation of the staleness of the data, analytically and sometimes even in closed form [1]. However, unless strong real-time requirements are imposed, it is rare to impose AoI minimization as the main objective. Very often, AoI is combined with other requirements related to more classic metrics of delay and/or throughput, as well as energy efficiency [4], [21], [22]. This might explain why this metric is still relatively unexplored by communication standards. Yet, we argue that with the advent of ultra-low latency and mission critical communications, in industrial, medical, and vehicular services [5], [10], the tide is going to turn.

In the majority of the investigations where AoI is introduced as a scheduling objective, the methodology is often borrowed from renewal processes and/or Markov chain, and the focus is on long-term performance and stable policies; therefore, an infinite-horizon average AoI is considered [6], [23]–[25].

However, finite-horizon optimization may possibly be more relevant. Even though these investigations of optimal steady-state policy undoubtedly have academic value, they hardly apply to the systems considered by networking standards and protocols such as industrial IoT systems and mission-critical communications, since, in these contexts, operations are often task driven and time bounded [7], [14], [26]. Processes such as predictive maintenance, quality control, or real-time actuation operate within pre-established time windows or production cycles. Thus, AoI optimization over a finite horizon would be more appropriate to align decisions with task-specific goals and deadlines, leading to more efficient and context-aware behavior than long-term or steady-state approaches [16].

Another aspect that may be the source of debate is the choice between average or peak AoI. Ideally, these two quantities describe different aspects of information freshness, the former being related to the average behavior, whereas the latter represents the worst case, i.e., it represents the most stale information possible, which is in turn contrasted with the requirement for the controller to make accurate decisions [15], [18]. Quite often, the dichotomy between these quantities is invoked to claim a supposedly different behavior. However, as we will show in the following, these two metrics behave quite similarly in the context of finite-horizon scheduling, which we believe is an important conclusion towards the definition of AoI-aware communication standards [27].

III. Methodology

Consider a transmitter and receiver exchanging data over a possibly noisy channel. We look for an AoI-optimal scheduler over a discrete time with a finite horizon equal to N time slots. We assume that the maximum number of transmission opportunities that the source can perform in the N slots is set to M. This implies that the source has a duty cycle constraint of $(M+1)/N$, which is in line with the energy and/or legal requirements of many communication standards.

Throughout our analysis, we consider some simplifying assumptions such as neglecting the propagation delay of the transmissions, assuming the feedback about the transmission outcome to be always correctly reported back at the transmitter, and new information to be always available at the transmitter's side when needed (a configuration called "generate at will") [23]. None of these assumptions is actually critical, as was shown in other papers that investigated them, so we do not involve them in the present study. For example, if there is a constant propagation delay, this is obviously outside of the optimization, and even if it is variable, AoI can still be optimized considering an approach in expectation, as argued in [26]. The relaxation of generation at will into sporadic update arrivals is the subject of [16], whereas imperfect feedback is analyzed by many papers, for example [18].

A. Average AoI minimization

We represent a system evolving over discrete time $t = 0, 1, \ldots, N$, and we let $a[t] \in \{0, 1, \ldots, N\}$ and $m[t] \in \{0, 1, \ldots, M\}$ denote the AoI in time slot t and the number of transmissions left, respectively. A minimization of average AoI can be obtained by considering the state of the system in the slot t as $x[t] = \big(a[t], m[t]\big)$, with $a[t]$ being the instantaneous AoI and $m[t] \in \{0, \ldots M\}$ the number of transmission opportunities left in the slot t. The system starts in state $x[0] = (0, M)$. According to the dynamic programming framework of [28], we can also consider the system control and noise processes as following. The former corresponds to a binary decision (0 or 1 corresponding to not transmit / transmit, respectively). The noise is included in the probability of an update being successful in resetting AoI. This is characterized by a single success probability being taken as $1 - \varepsilon$, where ε is the error rate, in case of independent errors. As per [29], this can also be extended to a Gilbert-Elliot-like correlated case by expanding the system state with the current channel condition and choosing two different values of the success probability depending on the current channel state. Further details on this model are given in

Thus, given the current $x[t]$, the system evolves into $x[t{+}1]$ that can be either $\big(a[t]{+}1, m[t]\big)$ if the source is silent in slot t. A transmission can be attempted instead only if $m[t] > 0$, in which case $m[t{+}1] = m[t] - 1$. Then, if the transmission is successful $a[t+1]$ is reset to 0, otherwise we still get $a[t{+}1] = a[t] + 1$. The optimal control policy $\mu_t\big(x[t]\big)$ is therefore defined as the choice that minimizes the expectation of a penalty equal to $a[t]$. A more in-depth discussion with full-fledged formalization is available in [14].

Here, we just discuss two immediate properties of μ_t. First of all, it can be obtained through backward induction; indeed, the optimal control policy μ_N in the last slot N is trivially to transmit, as long as $m[N] > 0$, since there is no use in saving transmission opportunities after the end of the time horizon. From this, the optimal choice at time $N{-}1$ is also immediately found by a binary comparison according to Bellman's optimality principle, and keeping propagating this procedure yields μ_t for any $t = 1, \ldots, N$. Second, as is intuitive and typical of AoI-based scheduling [30], μ_t has a threshold structure, i.e., if $\mu_t(a, m) = 1$ then $\mu_t(a', m') = 1$ for every $a' > a$ and $m' > m$.

B. Peak AoI minimization

To optimize Peak AoI instead, we need to expand the system state into $x[t] = \big(a[t], b[t], m[t]\big)$, with $b[t]$ being the highest AoI that has been experienced so far by the system, i.e., $b[t] = \min_{\tau=0,\ldots,t} a[\tau]$. The goal now becomes to minimize $\mathbb{E}\big[b[N]\big]$, the expected peak AoI at the end of the horizon.

Conveniently enough, this can still be performed as a dynamic program. The update law for $b[t]$ is $b[t{+}1] = \max(b[t], a[t{+}1])$. Also, a backward induction approach still works starting from a policy μ_N in the last slot that triggers a transmission if $m[t] > 0$. Compared to the previous optimization of the Average AoI, ties are more frequent (for example, when $b[t] \gg a[t]$, transmitting and staying idle have the same effect on the final Peak AoI); however, this can be solved through a proper rule for breaking ties. Specifically, if $m[t] > 0$, we chose transmission if $t > N{-}m$, and idling otherwise (this forces the last slot to use all the leftover transmission opportunities).

Another difference with the minimization of the average AoI is that the penalty minimization does not follow a combination of local cost and long-term expected penalty, since there is no per-step cost. Only the final value is important, and the penalty simply propagates the maximum AoI achieved in the horizon. However, the procedure still follows the same rationale. Finally, the optimal policy for minimum peak AoI has a larger space complexity, since it depends on three parameters (plus possibly the channel state) and therefore is $\mathcal{O}(N^3)$ compared to $\mathcal{O}(N^2)$ for minimum average AoI. However, we must point out that the optimal policy is precomputed and not evaluated at run-time. Also, it has, analogously to the average AoI minimization, a similar threshold structure that enables considerable simplifications in its memorization.

C. Channel Models: i.i.d. and Correlated Errors

In addition to the previous comparison, another extension of the standard AoI minimization studied in [14], [21], [23] is the channel correlation between the state of each slot. While it is immediate to represent independent and identically distributed (i.i.d) errors with a single parameter representing the error probability, an extended model where the channel state is correlated according to a Markov chain can also be conceived.

In this case, the error probability of the transmission is described by a Gilbert-Elliot-like process, whose state is included within the system state. In the simplest version, the channel has two states, namely "good" and "bad," and we take these to correspond to the error probability of 0 and 1 respectively. Thus, we can take the evolution of the channel state as following matrix $\mathbf{P} = \{p_{ij}\}_{i,j \in \{0,1\}}$, where $\varepsilon = p_{10}/(p_{01} + p_{10})$ is the steady state probability that the channel state is 1 (i.e., "bad"), and we need an additional parameter, usually taken as the average error burst length $B = 1/p_{10}$, to fully describe the channel. This model is widely used in communication systems for its simplicity, and is also adopted by some papers exploring AoI minimization under burst errors [19], [31], [32].

In [29], such a scenario is used to represent satellite communications, but the channel state refers to the current transmission, so that "bad" means that the satellite is out of sight and therefore no transmission can be performed. This is preventively known to the transmitter, which in this case can only adopt a policy $\mu_t = 0$. Conversely, in our scenario the channel state can evolve before the transmission is made, so that a transmission can be attempted even when the channel state is "bad," although this is generally less convenient because of the correlation.

IV. Results and Discussion

This section presents a comparative analysis between the policies minimizing the average or peak AoI, also showing the performance of the complementary non-optimized metric.

We consider two different channel models: an erasure channel with i.i.d errors and one with correlated errors according to the Gilbert-Elliot model described in III-C. We took the length of the horizon as $N{=}100$ slots, but the results displayed are normalized for ease of comparison and also to allow for an abstraction from this particular choice. We considered $M{=}4$ or $M{=}6$ transmission opportunities to happen within this frame. As a reference, in the absence of errors, the peak AoI should

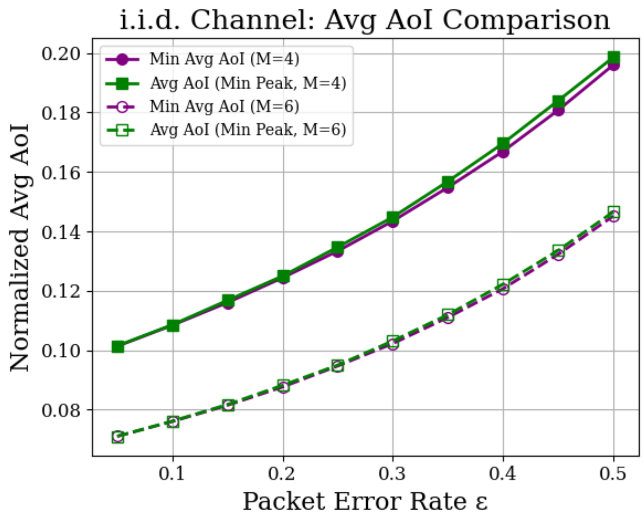

Fig. 1. Comparison of normalized average AoI vs. packet error rate ε under a channel model with i.i.d. errors. Results are averaged over 10^4 Monte Carlo runs, for an original horizon of N=100.

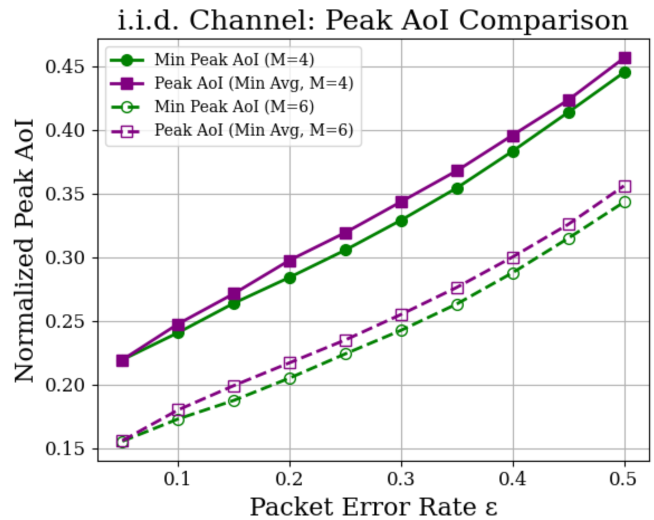

Fig. 2. Comparison of normalized peak AoI vs. packet error rate ε under a channel model with i.i.d. errors. Results are averaged over 10^4 Monte Carlo runs, for an original horizon of N=100.

be equal to $N/(M{+}1)$, and the normalized value is therefore just $1/(M{+}1)$, while the average AoI is clearly half of that.

We also remark that the values of M are generally and intentionally higher than what communication standards allow for. Generally speaking, nodes are allowed to transmit with a duty cycle of 1% or less [7]–[9], in which case there is even less slack for optimization, and the results that we will show in the following, basically proving a convergence of the metrics, hold true in an even stronger sense.

After finding the optimal policy through dynamic programming [28], we evaluated it using a Monte Carlo approach involving 10^4 realizations of the channel. The statistical confidence of the derived results (not shown to avoid overcrowding of the graphs) is fairly high and always above 99%.

In general, our findings indicate that although each policy excels in optimizing its own metric, the performance degradation of the alternative metric is frequently minimal. Moreover, this general conclusion is confirmed in both cases of i.i.d errors and channels with memory, as detailed in the following.

A. i.i.d Channel Comparison

We start by considering AoI minimization in the presence of i.i.d. channels, plotted against the error rate ϵ. We compare both AoI minimization policies of average and peak values.

As shown in Fig. 1, the average AoI values obtained by the two policies exhibit an extremely similar behavior, with the average AoI increasing as the error rate increased. Although the policy aimed at minimizing the average AoI achieves a slightly lower AoI, the gap is negligibly minimal. We can therefore infer that a minimization of peak AoI is also robust in terms of average AoI, achieving near-optimal performance.

Fig. 2 shows the reverse evaluation of the peak AoI for both policies, the one that actually minimizes the peak AoI and the one minimizing the average AoI instead. The performance gap here is slightly higher, especially as ε increases, yet the trends are once again very similar. This can be interpreted as a close-knit relationship between these two metrics and indeed proves that limiting the worst-case value also contains the average and vice versa.

B. Correlated Error Channel Comparison

We now analyze how average and peak AoI behave over a finite horizon in the presence of channels with correlated errors. To this end, we set the value of the error rate at ε=0.2 and we vary the burst length B chosen as our independent variable. We remark that even different choices of ε lead to qualitatively similar results.

Fig. 3 shows the average AoI as a function of B. When correlation grows and error bursts are longer, the average AoI increases. However, this applies to both policies and their results are fairly similar.

Fig. 4 shows instead the normalized peak AoI for different values of the average burst length B. Once again, the AoI-minimizing policy that considers the average AoI as the objective obtains slightly higher peak AoI values, but the curves are actually very close. Moreover, compared with the previous plots, the value of peak AoI is consistently almost flat, which is in line with what argued in [20] that the peak AoI is not affected by second-order statistics (in this case, the average burst length describes the channel correlation but the error rate is always the same).

V. Conclusions

We compared the minimization of average and peak AoI [15] for finite-horizon schedulers following dynamic programming recursions [14]. The key finding of our analysis is that the optimal transmission policies for average and peak AoI often converge to very similar performance; this seems to imply that

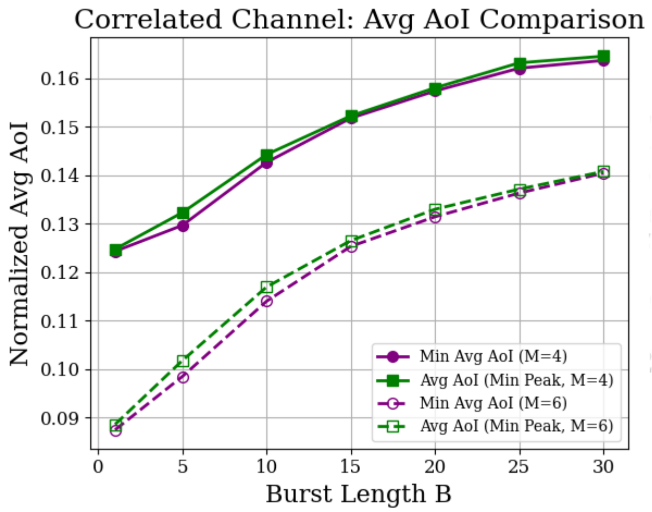

Fig. 3. Comparison of normalized average AoI vs. average burst length B under a channel model with correlated errors and average error rate ε=0.2. Results are averaged over 10^4 Monte Carlo runs, for an original horizon of N=100.

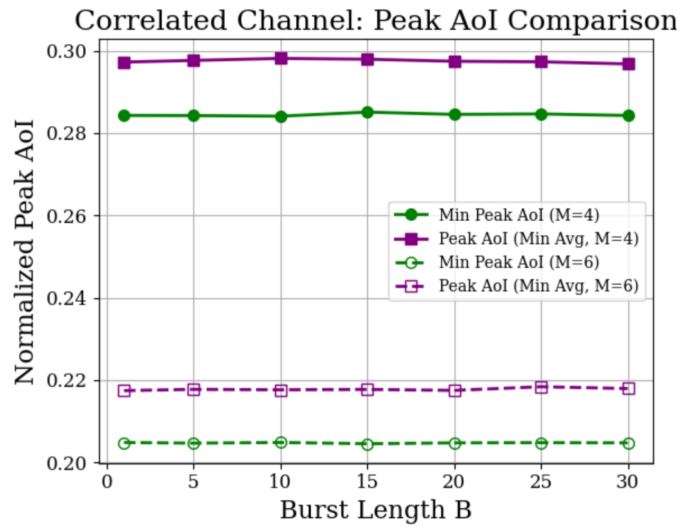

Fig. 4. Comparison of normalized peak AoI vs. average burst length B under a channel model with correlated errors and average error rate ε=0.2. Results are averaged over 10^4 Monte Carlo runs, for an original horizon of N=100.

simplified and unified policies may suffice for many finite-horizon applications if a general requirement of information freshness is sought.

In addition, correlated errors cause an increase in the average AoI, as it becomes more likely to miss subsequent updates. However, the peak AoI is not affected by this trend, being only affected by first-order statistics [20]. In any event, both of these trends occur irrespectively of the exact metric of choice being the average or peak AoI.

These results point to several promising paths. The analysis can be extended with multiple coexisting sources, which possibly involves considering medium access and coordination through game theory [33]. Furthermore, incorporating other aspects, such as energy constraints, would better reflect resource-limited IoT devices. All of these considerations may find application likely soon in the context of communication standards for next-generation networks, in light of the increasing need for real-time services based on information freshness.

References

[1] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, “Age of information: An introduction and survey,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, 2021.

[2] D. Sinha and R. Roy, “Scheduling status update for optimizing age of information in the context of industrial cyber-physical system,” *IEEE Access*, vol. 7, pp. 95 677–95 695, 2019.

[3] M. Vaezi, K. Noroozi, T. D. Todd, D. Zhao, and G. Karakostas, “Digital twin placement for minimum application request delay with data age targets,” *IEEE Internet Things J.*, vol. 10, no. 13, pp. 11 547–11 557, 2023.

[4] I. Kadota, A. Sinha, and E. Modiano, “Scheduling algorithms for optimizing age of information in wireless networks with throughput constraints,” *IEEE/ACM Trans. Netw.*, vol. 27, no. 4, pp. 1359–1372, 2019.

[5] I. Kahraman, A. Köse, M. Koca, and E. Anarim, “Age of information in Internet of things: A survey,” *IEEE Internet Things J.*, vol. 11, no. 6, pp. 9896–9914, 2023.

[6] Y.-P. Hsu, E. Modiano, and L. Duan, “Scheduling algorithms for minimizing age of information in wireless broadcast networks with random arrivals,” *IEEE Trans. Mobile Comput.*, vol. 19, no. 12, pp. 2903–2915, 2019.

[7] A. Rolich, I. Turcanu, A. Vinel, and A. Baiocchi, “Understanding the impact of persistence and propagation on the age of information of broadcast traffic in 5G NR-V2X sidelink communications,” *Comp. Netw.*, vol. 248, p. 110503, 2024.

[8] LoRa Alliance, “LoRaWAN 1.1 specification,” vol. 11, 2018.

[9] F. Fernández, M. Zverev, P. Garrido, J. R. Juárez, J. Bilbao, and R. Agüero, “Even lower latency in IIoT: Evaluation of QUIC in industrial IoT scenarios,” *Sensors*, vol. 21, no. 17, p. 5737, 2021.

[10] T. Stüber, L. Osswald, S. Lindner, and M. Menth, “A survey of scheduling algorithms for the time-aware shaper in time-sensitive networking (TSN),” *IEEE Access*, vol. 11, pp. 61 192–61 233, 2023.

[11] E. Uysal, O. Kaya, A. Ephremides, J. Gross, M. Codreanu, P. Popovski, M. Assaad, G. Liva, A. Munari, B. Soret *et al.*, “Semantic communications in networked systems: A data significance perspective,” *IEEE Netw.*, vol. 36, no. 4, pp. 233–240, 2022.

[12] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, “Minimizing age of information in vehicular networks,” in *Proc. IEEE SECON*, 2011, pp. 350–358.

[13] M. K. Abdel-Aziz, S. Samarakoon, C.-F. Liu, M. Bennis, and W. Saad, “Optimized age of information tail for ultra-reliable low-latency communications in vehicular networks,” *IEEE Trans. Commun.*, vol. 68, no. 3, pp. 1911–1924, 2019.

[14] A. Buratto, H. Tuwei, and L. Badia, “Optimizing sensor data transmission in collaborative multi-sensor environments,” in *Proc. IEEE Int. Conf. Commun. Netw. Satellite (COMNETSAT)*, 2023, pp. 635–639.

[15] T. Zhang, S. Chen, Z. Chen, Z. Tian, Y. Jia, M. Wang, and D. O. Wu, “AoI and PAoI in the IoT-based multisource status update system: Violation probabilities and optimal arrival rate allocation,” *IEEE Internet Things J.*, vol. 10, no. 23, pp. 20 617–20 632, 2023.

[16] L. Badia and A. Munari, “Exogenous update scheduling in the industrial Internet of things for minimal age of information,” *IEEE Trans. Ind. Informat.*, vol. 21, no. 2, pp. 1210–1219, 2025.

[17] J. P. Champati, R. R. Avula, T. J. Oechtering, and J. Gross, “Minimum achievable peak age of information under service preemptions and request delay,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1365–1379, 2021.

[18] F. Chiariotti, A. Munari, L. Badia, and P. Popovski, “Peak age of incorrect information of reactive ALOHA reporting under imperfect feedback,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2025, pp. 1–6.

[19] L. Badia, "Analysis of age of information under SR ARQ," *IEEE Commun. Lett.*, vol. 27, no. 9, pp. 2308–2312, 2023.
[20] P. Zou, O. Ozel, and S. Subramaniam, "Waiting before serving: A companion to packet management in status update systems," *IEEE Trans. Inf. Theory*, vol. 66, no. 6, pp. 3864–3877, 2019.
[21] S. Feng and J. Yang, "Age of information minimization for an energy harvesting source with updating erasures: Without and with feedback," *IEEE Trans. Commun.*, vol. 69, no. 8, pp. 5091–5105, 2021.
[22] E. T. Ceran, D. Gündüz, and A. György, "A reinforcement learning approach to age of information in multi-user networks with HARQ," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1412–1426, 2021.
[23] A. Javani, M. Zorgui, and Z. Wang, "On the age of information in erasure channels with feedback," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2020, pp. 1–6.
[24] C. Kam, S. Kompella, and A. Ephremides, "Age of information under random updates," in *Proc. IEEE Int. Symp. Inf. Th. (ISIT)*, 2013, pp. 66–70.
[25] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2637–2650, 2018.
[26] L. Badia, A. Zancanaro, G. Cisotto, and A. Munari, "Status update scheduling in remote sensing under variable activation and propagation delays," *Ad Hoc Netw.*, vol. 163, p. 103583, 2024.
[27] A. Buratto, A. Munari, and L. Badia, "Strategic backoff of slotted ALOHA for minimal age of information," *IEEE Commun. Lett.*, vol. 29, no. 1, pp. 155–159, 2025.
[28] D. Bertsekas, *Dynamic programming and optimal control: Volume I*. Athena scientific, 2012, vol. 4.
[29] L. Badia and A. Munari, "Satellite intermittent connectivity and its impact on age of information for finite horizon scheduling," in *Proc. Adv. Satellite Multimedia Syst. Conf. Signal Process. Space Commun. Wkshp (ASMS/SPSC)*, 2025, pp. 1–8.
[30] E. Fountoulakis, T. Charalambous, A. Ephremides, and N. Pappas, "Scheduling policies for AoI minimization with timely throughput constraints," *IEEE Trans. Commun.*, vol. 71, no. 7, pp. 3905–3917, 2023.
[31] Z. Wang and G. B. Giannakis, "A simple and general parameterization quantifying performance in fading channels," *IEEE Trans. Commun.*, vol. 51, no. 8, pp. 1389–1398, 2003.
[32] M. Rossi, L. Badia, and M. Zorzi, "SR ARQ delay statistics on N-state Markov channels with non-instantaneous feedback," *IEEE Trans. Wireless Commun.*, vol. 5, no. 6, pp. 1526–1536, 2006.
[33] L. Badia, "Age of information from two strategic sources analyzed via game theory," in *Proc. IEEE Int. Worksh. Comp. Aided Model. Design Commun. Links Netw. (CAMAD)*, 2021, pp. 1–6.