

Spatio-Temporal Information Freshness for Remote Source Monitoring in IoT Systems

Andrea Munari, Federico Chiariotti, Leonardo Badia, Petar Popovski

Abstract—The widespread adoption of age of information (AoI) as a meaningful and analytically tractable information freshness metric has led to a wide body of work on the timing performance of Internet of things (IoT) systems. However, the spatial correlation inherent to environmental monitoring has been mostly neglected in the recent literature, due to the significant modeling complexity it introduces. In this work, we address this gap by presenting a model of spatio-temporal information freshness, considering the conditional entropy of the system state in a remote monitoring scenario, such as a low-orbit satellite collecting information from a wide geographical area. Our analytical results show that purely age-oriented schemes tend to select an overly broad communication range, leading to inaccurate estimates and energy inefficiency, both of which can be mitigated by adopting a spatio-temporal approach.

I. INTRODUCTION

Over the past decade, billions of IoT sensors have been deployed to enable a plethora of different applications, from industrial automation to environmental monitoring [1]. The unique challenges of dealing with massive numbers of low-power, low-throughput devices have reshaped medium access design [2], as well as the performance metrics, which fundamentally differ from those used for human communication.

Rather than aiming at minimizing the latency of every single packet, AoI [3] has emerged as a compact and useful metric for IoT scenarios. The metric measures the freshness of the last update that was correctly received, offering a better proxy for the alignment between the estimate available to the receiver and the real state of the environment. Since its introduction, AoI and its optimization have been the subject of a vast body of scientific literature [4], as well as several extensions.

While the timing aspects of information freshness have been thoroughly investigated, the spatial dimension is so far mostly unexplored, or dealt with only as part of the medium access model [5]: research that considers both dimensions, exploiting both age-related notions and spatial correlations in the physical process measured by the sensors, is still relatively sparse. The first work to explicitly model correlated sources [6] used beliefs at the receiver to compute a spatio-temporal AoI metric and optimize receiver-driven scheduling




Fig. 1. A schematic of the considered application scenario.

of sensor observations. Another model, which starts from similar considerations, was presented in [7]: in this case, simpler access strategies are considered, which allow for a more complete freshness analysis.

Other possible models adapt the definition of AoI without explicitly considering the receiver's beliefs: it is possible to reset the age to a value that reflects the correlation of the received measurement with the intended target, artificially increasing it to account for imperfect measurements, or to adapt the age function [8], considering a distance-based weighting function so that age increases faster for less reliable measurements [9]. Sensor correlation is also a critical factor when considering more complex environment models such as digital twins [10].

In this work, we consider an application scenario such as the one depicted in Fig. 1: the physical process that the network aims at measuring, represented by the tree, is localized in space, and only sensors placed very close to it may measure it exactly. However, we cannot directly control sensor placement, and energy efficiency considerations prevent us from only considering very close sensors, as it would quickly deplete their batteries. The correlation of each sensor's observations with the real state of the process depends on the distance between them, represented by the shaded areas in the figure. Measurements from sensors closer to the tree are highly correlated to its state, while sensors that are placed too far from it can only make very unreliable observations. The receiver, represented by the satellite, can collect observation within a controllable radius and track the process. In this case, pure AoI is insufficient to represent the freshness of the information available to the satellite: similarly to the model from [6], we need to consider both the freshness of updates and their uncertainty to account for unreliable measurements.

Unlike [6], which studies a centralized scheduling scheme in which the receiver perfectly knows the placement of the sensors, we consider a slotted ALOHA access protocol, which

A. Munari (andrea.munari@dlr.de) is with the Inst. of Comm. and Navigation, German Aerospace Center (DLR), Germany. F. Chiariotti (federico.chiariotti@unipd.it) and L. Badia (leonardo.badia@unipd.it) are with the Dept. of Information Eng., Univ. of Padova (Italy), P. Popovski (petarp@es.aau.dk) is with the Dept. of Electronic Systems, Aalborg University (Denmark). This work was supported in part by the Velux Foundation, Denmark, through the Villum Investigator Grant WATER, nr. 37793. A. Munari acknowledges the financial support by the Federal Ministry of Research, Technology and Space of Germany in the programme of "Souverän. Digital. Vernetzt." Joint project 6G-RIC, project id. number: 16KISK022.

is more realistic for IoT systems, and assume that nodes do not have any geolocation information. In this case, the uncertainty over the state of the source process can be characterized by the conditional entropy at the receiver. In this work, we consider a *forgetful* receiver that only tracks the age and the last received value. This makes the analysis tractable, as well as allowing for implementation on constrained hardware. We leave more complex setups for future work.

We provide a novel closed-form analysis of the conditional entropy, which reveals the importance of including a spatial component in information freshness problems: the minimum entropy solution is obtained by setting a much narrower radius than the AoI-optimal one, reducing energy consumption by limiting the number of polled sensors and improving the accuracy of the remote estimation. The optimization of the collection radius also has significant practical implications, offering hints on system dimensioning.

Notation: We denote vectors and matrices in lower- and upper-case boldface, respectively, e.g., $\mathbf{a}$, $\mathbf{A}$. In particular, a vector denoting the position of a point on $\mathbb{R}^2$ is indicated as $\mathbf{x}$. A discrete random variable (r.v.) and its realization use upper- and lower-case, respectively, e.g., X and x . Its probability mass function is denoted as $p_X(x) = P(X=x)$, and the subscript is omitted for brevity whenever it can be inferred from the context. The notation is extended straightforwardly to conditional distributions. We express the indicator function as $\mathbb{1}(x, y)$, taking value 1 if $x = y$ and 0 otherwise, and we denote the complementary of a binary value $y \in \{0, 1\}$ as $\bar{y}$.

II. SYSTEM MODEL

Throughout our study, we focus on a set of IoT nodes (devices), uniformly distributed over the circular coverage area $\mathcal{A} \subseteq \mathbb{R}^2$ of a receiver. Without loss of generality, $\mathcal{A}$ is centered at the origin of the plane, and has radius R_m . For a density ρ [nodes/unit area], we consider $m = \rho \pi R_m^2$ devices to populate the system, observing a physical process of interest and reporting updates to the receiver. The process is modeled as a discrete-time, two-state Markov chain X_n , taking values in $\mathcal{X} = \{0, 1\}$. The one-step transition matrix for the chain is

$$\mathbf{A} = \begin{pmatrix} 1-q & q \\ \eta q & 1-\eta q \end{pmatrix}$$

where $\eta > 0$ is referred to as *asymmetry factor*.

Devices share a common wireless channel towards the receiver by means of a slotted ALOHA access policy [11]. Specifically, time is divided in slots of equal duration, corresponding to the time steps of the Markov process, and at each slot $n \in \mathbb{N}$: i) X_n transitions following $\mathbf{A}$, and ii) every device independently decides with probability ζ whether to observe the process and transmit the obtained value. A sent packet is erased with probability ε , bringing no power contribution to the receiver, or arrives unfaded with probability $1 - \varepsilon$.¹

¹The considered model captures setups in which packets coming from different nodes are equally likely to be retrieved. For instance, this is the case under power control, or in settings where the relative displacement of nodes is not predominant on the overall path-loss, e.g., for a LEO satellite.

Resorting to the well-established collision channel model, we assume that a slot where two or more unerased messages reach the receiver does not allow retrieval of information, whereas a singleton slot leads to successful decoding [11]. Accordingly, each slot sees delivery of a message with probability

$$p_s = m \zeta (1 - \varepsilon) \cdot [1 - \zeta (1 - \varepsilon)]^{m-1}.$$

Due to the displacement of devices within $\mathcal{A}$, the reliability of the observations produced by a node depends on its position. In particular, we assume that devices closer to the origin of the plane provide more precise readings, whereas the accuracy of messages coming from farther away progressively reduces, as the longer the information about the event propagates physically, the higher the probability that it may be distorted. Such a setting captures, for instance, the case of an application tracking the value of a physical quantity at the specific coordinates of interest (plane origin), or of a monitoring system that has to detect an event (e.g., presence/absence of an object) in the corresponding surroundings. To model this, we partition the coverage area $\mathcal{A}$ in K regions:

$$\mathcal{A}_i := \{\mathbf{x} \in \mathcal{A} : iR \leq \|\mathbf{x}\|_2 < (i+1)R\}$$

for $i \in \mathcal{K} = \{0, \dots, K-1\}$, and where $R := R_m/K$. Whenever a device in $\mathcal{A}_i$ transmits an update, its message contains the correct value of X_n with probability $\lambda(i)$, or the erroneous reading $\bar{X}_n$ with probability $1 - \lambda(i)$. In the remainder of our discussion we set

$$\lambda(i) = (1 + iR)^{-\alpha} \quad (1)$$

with $\alpha > 0$, although the analysis we present can be adapted straightforwardly to different reliability functions.

At the receiver side, we assume that no knowledge can be gathered on the area out of which an incoming message was generated, e.g., due to location-unaware devices, or to the fact that no geographical information is added to the packet payload. Moreover, we consider the receiver to be *forgetful*, i.e., that it may estimate the state of the process of interest solely based on the last obtained reading. This solution is inspired by practically relevant IoT setups, epitomizing the common behavior of an application layer which is fed with the payload of incoming messages, and does not lean on any other cross-layer information exchange or memory. More intelligent models, which also require a higher computational power, will be considered as future work. Following this approach, the receiver at time n has knowledge about: (i) the time elapsed since the last reception, i.e., the age of information Δ_n [4]; and (ii) the content of the last retrieved message, which we denote by Y_n . Without loss of generality, the reception of a message resets the AoI value to $\Delta_n = 0$.

The uncertainty on the current state of the process X_n , conditioned on the available information, is then naturally captured by the conditional entropy

$$h(y_n, \delta_n) := H(X_n | Y_n = y_n, \Delta_n = \delta_n).$$

The quantity is relevant from an operational standpoint, as it may serve as basis to properly gauge the current knowledge or




Fig. 2. Examples of time evolution of $h(y_n, \delta_n)$ when tracking a symmetric source. The horizontal line reports the value of the stationary source entropy, $H(X) = -\pi_0 \log_2 \pi_0 - \pi_1 \log_2 \pi_1$. The plot was generated for $\alpha=0.02$, $R=10$, $q=0.005$, $\rho=0.05$, $\zeta=5 \cdot 10^{-4}$, $\varepsilon=0.1$.

make actuation decisions, and is akin to the age of uncertainty introduced in [12]. Two examples of how $h(y_n, \delta_n)$ may evolve over time are reported in Fig. 2. In the first case, (a), a smaller radius of $R_m = 6R$ was considered, whereas the rightmost plot, (b), was obtained for a wider coverage of $R_m = 15R$. The figure hints at a key tradeoff. On the one hand, selecting a smaller radius leads to more sporadic update deliveries (corresponding to resets of $h(y_n, \delta_n)$), and the uncertainty tends to grow for longer periods of time.

Whenever the receiver does not retrieve messages for some time (central part of plot (a)), its uncertainty tends to converge to the stationary entropy of the source, i.e., $H(X) = -\pi_0 \log_2 \pi_0 - \pi_1 \log_2 \pi_1$. This effect is much more mitigated in figure (b), as the presence of more nodes in the system results in an improvement of the frequency of message delivery, with a lower AoI and an uncertainty that is more often brought back to its minimum. On the other hand, such reset value gets higher as R_m increases. Indeed, incoming messages are more likely to be generated by devices which are farther away from the origin, containing less reliable readings as per (1), and leaving the receiver with a higher residual uncertainty. The critical balance of these two factors will determine system performance, and its analytical characterization represents the core contribution of this paper. We will provide in Sec. III a closed-form derivation of the conditional entropy

$$H(X_n | Y_n, \Delta_n) = \sum_{y_n, \delta_n} h(y_n, \delta_n) p(y_n, \delta_n) \quad (2)$$

as well as the cumulative distribution function of $h(y_n, \delta_n)$.

III. ANALYSIS

Let us start by considering a slot n in which the receiver successfully retrieves a message, i.e., the AoI is reset to $\Delta_n=0$. Applying Bayes' rule, the distribution of the state of the tracked process conditioned on receiving a reading with value Y_n can be expressed as

$$p_{X_n | Y_n, \Delta_n}(x_n | y_n, 0) = \frac{p_{Y_n | X_n, \Delta_n}(y_n | x_n, 0) \pi_{x_n}}{\sum_{x'_n \in \mathcal{X}} p_{Y_n | X_n, \Delta_n}(y_n | x'_n, 0) \pi_{x'_n}} \quad (3)$$

where π_{x_n} is the stationary distribution of X_n , readily obtained as $\pi_0 = \eta/(1+\eta)$, $\pi_1 = 1/(1+\eta)$, and we have used the fact that event of successfully receiving a packet is independent of X_n for the transmission policies under study.

In turn, $p(y_n | x_n, 0)$ can be conveniently expressed in terms of the r.v. $D_n \in \mathcal{K}$, capturing the region the message originated from. Leaning on the law of total probability, we have

$$p(y_n | x_n, 0) = \sum_{d_n=0}^{K-1} p(y_n | x_n, 0, d_n) p(d_n | x_n, 0). \quad (4)$$

Within (4), the term $p(d_n | x_n, 0)$ accounts for the probability of a received message to be generated from region d_n . Observing that the devices are uniformly distributed within $\mathcal{A}$, and transmit independently of their location as well as of the state of the process, the quantity can be computed as

$$p(d_n | x_n, 0) = \int_{d_n R}^{(d_n+1)R} \frac{2r}{R_m^2} dr = \frac{2d_n + 1}{K^2}. \quad (5)$$

The first factor within the summation in (4) captures the probability of a decoded message containing value y_n , conditioned on the actual state of the process and on the reliability of the sender. Accordingly, we can write in a compact form

$$p(y_n | x_n, 0, d_n) = \lambda(d_n) \mathbb{1}(x_n, y_n) + (1-\lambda(d_n)) \mathbb{1}(\bar{x}_n, y_n). \quad (6)$$

For the power-law in (1), combining (6) and (5) we obtain

$$p(y_n | x_n, 0) = \sum_{d_n=0}^{K-1} \frac{2d_n + 1}{K^2(1 + d_n R)^\alpha}, \quad \text{if } y_n = x_n \quad (7)$$

and its complementary value for the case $y_n \neq x_n$. Plugging (7) into (3) thus provides the statistics available at the receiver on the state of the process of interest upon reception of a packet. The result, in turn, allows us to derive the general probability $p(x_n | y_n, \delta_n)$ for any $\delta_n > 0$, obtained as the δ_n -step evolution of the transition matrix $\mathbf{A}$ with initial distribution $p_{X_n | Y_n, \Delta_n}(0 | y_n, 0)$ and $p_{X_n | Y_n, \Delta_n}(1 | y_n, 0)$.

The uncertainty that the forgetful receiver has on the current state of the process given the currently available knowledge can finally be computed by definition as

$$h(y_n, \delta_n) = - \sum_{x_n \in \mathcal{X}} p(x_n | y_n, \delta_n) \log_2 p(x_n | y_n, \delta_n). \quad (8)$$

To conclude, we also derive the conditional entropy $H(X_n | Y_n, \Delta_n)$. From (2), this requires the PMF $p(y_n, \delta_n)$, capturing, for a generic time slot n , the joint probability of having decoded the last message δ_n slots ago, and of the received reading being y_n . Observing that packet receptions are i.i.d. across slots with probability p_s , the current AoI value is independent of the content of the incoming message, as each node's transmission probability does not depend on its age or position. Accordingly, $p(y_n, \delta_n) = p(y_n)p(\delta_n)$. Moreover, $p(y_n)$ can simply be computed from (4) as

$$p(y_n) = \sum_{x_n \in \mathcal{X}} p(y_n | x_n, 0) \pi_{x_n}. \quad (9)$$




Fig. 3. Conditional entropy $H(X_n | Y_n, \Delta_n)$ vs. coverage radius R_m , when tracking a symmetric (dashed line) or an asymmetric (solid line) process. The markers report the outcomes of simulation results. The vertical line indicates the coverage radius that would minimize AoI.

In turn, the receiver experiences an AoI $\delta_n \geq 0$ at the start of slot n if the last success was followed by $\delta_n - 1$ failures:

$$p(\delta_n) = p_s(1 - p_s)^{\delta_n - 1}. \quad (10)$$

Combining (9)-(10) with (8) provides the conditional entropy.

Remark: The approach followed in the analysis above considers an i.i.d. distribution for the random variable D_n across slots. This implicitly corresponds to independently re-drawing the positions of all nodes at each time unit, and thus introduces an approximation of the behavior of a practical topology in which the locations of the devices are static. The validity of this assumption will be discussed in Sec. IV.

IV. RESULTS AND DISCUSSION

To gauge the impact of spatio-temporal information freshness on the forgetful receiver, we focus on a population of devices spread over $\mathcal{A}$ with density $\rho = 5 \cdot 10^{-2}$ [nodes/m²], transmitting at every slot with probability $\zeta = 10^{-4}$. A sent packet is erased with probability $\varepsilon = 0.1$, and the tracked Markov process X_n transitions from state 0 to 1 with probability $q = 5 \cdot 10^{-3}$. Unless otherwise specified, we assume that nodes produce readings following the reliability law in (1), where $\alpha = 0.02$, and $R = 10$ m.

We start our discussion by reporting in Fig. 3 the conditional entropy $H(X_n | Y_n, \Delta_n)$ against the coverage radius R_m . The dashed line reports the case of monitoring of a symmetric process ($\eta = 1$), whereas the solid curve was obtained for an asymmetric case with $\eta = 5$. Circle markers show the results of detailed network simulations, aimed at verifying the approximation introduced in our analysis (see Remark 1). Specifically, for each simulation, the positions of the m nodes were uniformly drawn within $\mathcal{A}$, and their location was kept fixed for the whole duration of the run (10^6 slots). Throughout the run, the Markov process evolved according to $\mathbf{A}$, and the time evolution of $h(y_n, \delta_n)$ was computed via (8) to derive its average value. Each of the reported marked points shows in turn the average of 20 independent topologies. The excellent match between analysis and simulations under all conditions confirms the tightness of the modeling assumptions.




Fig. 4. (a) coverage range R_m^* minimizing the conditional entropy of the forgetful receiver; (b) corresponding minimum value of $H(X_n | Y_n, \Delta_n)$.

More interestingly, Fig. 3 reveals the existence of an optimal coverage range to minimize $H(X_n | Y_n, \Delta_n)$. When coverage is too small, updates are sporadic as few nodes fall within $\mathcal{A}$. In this case, even if incoming messages are highly reliable in view of the proximity of active devices to the origin, the receiver's uncertainty is dominated by the lack of timely (fresh) updates. On the other hand, if the coverage radius is increased beyond a critical point R_m^* , the conditional entropy grows once more (rightmost region of the plot). In such conditions, even if messages are received more often, their reliability decreases sharply due to the higher probability of obtaining readings from nodes farther away from the origin, leading to a higher uncertainty. Striking a proper balance between the temporal and spatial components of information freshness is thus paramount, and shall represent a key criterion in the design of the system. From this standpoint, it is relevant to remark that the optimal radius R_m^* can be significantly smaller than the one that would be obtained considering AoI only. This is reported for reference by the vertical dash-dotted line in the figure. For the generate-at-will case under study, such configuration is obtained by operating the system at maximum throughput [13], [14], i.e., a channel load of 1 [pkt/slot], with a corresponding radius $(\pi\rho\zeta\varepsilon)^{-1/2}$.

Fig. 3 also highlights a difference in the optimal coverage when tracking symmetric and asymmetric sources, with the latter case presenting lower conditional entropy and a larger R_m^* . As η increases, the tracked process spends more time in state 0, and the uncertainty on its current conditions tends to converge more quickly to a lower value even in the absence of incoming updates. This effect reduces the impact of the spatial component of information freshness, as a low level of reliability of incoming messages can be compensated by the knowledge of the biased behavior of X_n .

To further delve into this aspect, we report in Fig. 4 the optimal coverage radius R_m^* , (a), and the corresponding minimum conditional entropy $H^*(X_n | Y_n, \Delta_n)$, (b), against the asymmetry factor. The optimal radius sharply increases as the asymmetry grows, and eventually, R_m^* tends to converge to the value minimizing AoI, i.e., $(\pi\rho\zeta\varepsilon)^{-1/2}$, as the temporal

component of freshness dominates. A corresponding sharp decrease in the attainable average uncertainty is experienced as η grows, obtained by admitting more nodes in the system. The plots report trends for two different values of the exponent α , determining the reliability of incoming messages as per (1). It is interesting to note that, for larger values of α , the optimal coverage shrinks, as an effect of the quicker degradation with distance of the quality of the updates. In this setting, trading off messages' timeliness for an increased reliability pays off. For instance, when $\eta = 5$ (same setup considered in Fig. 3), the optimal coverage radius reduces to less than one half of what would be optimal in terms of when $\alpha = 0.06$.

We conclude our discussion by reporting in Fig. 5 the cumulative distribution function of the uncertainty at the receiver at a generic point in time, conditioned on the last received message and its AoI. The metric can be obtained from (8), (9)-(10) as

$$P(h(y_n, \delta_n) \leq w) = \sum_{\substack{(y_n, \delta_n) \text{ s.t.} \\ h(y_n, \delta_n) \leq w}} p(y_n, \delta_n).$$

In the plot, we focus on a symmetric process X_n , and consider two different coverage radii: a relatively small one (25 [m], dashed lines) and a larger one (125 [m], solid lines). Curves without markers were obtained for $\alpha = 0.02$, whereas the ones with markers denote the behavior with $\alpha = 0.06$. Consider first the former case. When a narrower $\mathcal{A}$ is covered, the receiver's uncertainty is reset to a smaller value whenever a message is successfully decoded, by virtue of the more reliable readings being collected. As a result, the dashed line reaches values of $h(y_n, \delta_n)$ close to 0.3 [bit] in the leftmost part of the figure, while the uncertainty never falls below 0.4 [bit] for $R_m = 125$ [m]. On the other hand, the availability of more transmitting devices leads to much more frequent deliveries. For the considered settings, with $R_m = 125$ [m], we get $p_s \simeq 0.1$, and a new reading is obtained on average every ~ 9 slots. This results in a distribution of $h(y_n, \delta_n)$ which is much more concentrated around the value of uncertainty that follows a reset, and leads to a lower average conditional entropy (see Fig. 3). Similar trends are observed for $\alpha = 0.06$, with a shift towards higher uncertainty levels, due to the lower reliability of incoming messages.

V. CONCLUSIONS

This work provides an initial analysis of the uncertainty of an IoT receiver over a spatio-temporal sensing task: along with the freshness dimension, which is traditionally accounted for through the AoI, our analysis also includes a spatial component, assigning a lower reliability to updates from sensors farther from the center of a monitored area. Our model of spatio-temporal freshness extends and complements results from the recent literature, which considered centralized models, by considering the optimization of a simple random access scheme, drawing some basic design guidelines.

Naturally, this is only a first step in the characterization of spatio-temporal freshness: the consideration of more complex




Fig. 5. Cumulative distribution function of the uncertainty $h(y_n, \delta_n)$ at the receiver. Results obtained for a symmetric process X_n , and for $\alpha = 0.02$.

estimators at the receiver, such as a hidden Markov model, would be an interesting avenue for further research, as would considering explicit information about each transmission's geolocation. Additionally, the study of different and more descriptive models for the source process and observation noise, such as the Ornstein-Uhlenbeck model, or even structural correlation models not purely tied to distance, would broaden the applicability of the results and provide further insights on the interplay of space, time, and value in remote monitoring.

REFERENCES

- [1] Ericsson, "Mobility report – November 2024," Ericsson, Tech. Rep., 2024.
- [2] F. Guo, F. R. Yu, H. Zhang, X. Li, H. Ji, and V. C. Leung, "Enabling massive IoT toward 6G: A comprehensive survey," *IEEE Internet of Things J.*, vol. 8, no. 15, pp. 11 891–11 915, 2021.
- [3] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Proc. IEEE SECON*, June 2011.
- [4] R. Yates, Y. Sun, D. Brown, S. 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, May 2021.
- [5] H. H. Yang, A. Arafa, T. Q. Quek, and H. V. Poor, "Spatiotemporal analysis for age of information in random access networks under last-come first-serve with replacement protocol," *IEEE Trans. Wireless Commun.*, vol. 21, no. 4, pp. 2813–2829, 2021.
- [6] A. E. Kalør and P. Popovski, "Timely monitoring of dynamic sources with observations from multiple wireless sensors," *IEEE/ACM Trans. Netw.*, vol. 31, no. 3, pp. 1263–1276, 2022.
- [7] A. Zancanaro, G. Cistotto, and L. Badia, "Modeling value of information in remote sensing from correlated sources," *Computer Communications*, vol. 203, pp. 289–297, 2023.
- [8] J. Tong, L. Fu, and Z. Han, "Age-of-information oriented scheduling for multichannel IoT systems with correlated sources," *IEEE Trans. Wireless Commun.*, vol. 21, no. 11, pp. 9775–9790, 2022.
- [9] M. Fidler, F. Gallistl, J. P. Champati, and J. Widmer, "2D-AoI: Age-of-information of distributed sensors for spatio-temporal processes," *arXiv preprint arXiv:2412.12789*, 2024.
- [10] P. Bellavista, N. Bricocchi, M. Fogli, C. Giannelli, M. Mamei, and M. Picone, "ODTE: A metric for digital twin entanglement," *IEEE Open J. Commun. Soc.*, 2024.
- [11] N. Abramson, "The throughput of packet broadcasting channels," *IEEE Trans. Commun.*, vol. COM-25, no. 1, pp. 117–128, 1977.
- [12] G. Chen, S.-C. Liew, and Y. Shao, "Uncertainty-of-information scheduling: A restless multiarmed bandit framework," *IEEE Trans. Inf. Theory*, vol. 68, no. 9, pp. 6151–6173, 2022.
- [13] R. Yates and S. Kaul, "Status updates over unreliable multiaccess channels," in *Proc. IEEE ISIT*, 2017.
- [14] A. Munari, "Modern random access: an age of information perspective on irregular repetition slotted ALOHA," *IEEE Trans. Commun.*, vol. 69, no. 6, pp. 3572–3585, 2021.