

Goal-Oriented Medium Access with Distributed Belief Processing

Federico Chiariotti, *Senior Member, IEEE*, Andrea Munari, *Senior Member, IEEE*,
 Leonardo Badia, *Senior Member, IEEE*, and Petar Popovski, *Fellow, IEEE*

Abstract—Goal-oriented communication entails the timely transmission of updates related to a specific goal defined by the application. In a distributed setup with multiple sensors, each individual sensor knows its own observation and can determine its freshness, as measured by Age of Incorrect Information (AoII). This local knowledge is suited for distributed medium access, where the transmission strategies have to deal with collisions. We present Dynamic Epistemic Logic for Tracking Anomalies (DELTA), a medium access protocol that limits collisions and minimizes AoII in anomaly reporting over dense networks. Each sensor knows its own AoII, while it can compute the belief about the AoII for all other sensors, *based on their Age of Information (AoI)*, which is inferred from the acknowledgments. This results in a goal-oriented approach based on dynamic epistemic logic emerging from public information. We analyze the resulting DELTA protocol both from a theoretical standpoint and with Monte Carlo simulations, showing that it is significantly more efficient and robust than classical random access, while outperforming state-of-the-art scheduled schemes by at least 30%, even with imperfect feedback.

Index Terms—Goal-oriented communication; age of incorrect information; dynamic epistemic logic; medium access control.

I. INTRODUCTION

Goal-oriented communication is a new paradigm that aims at overcoming the limits of traditional communication systems by considering the *meaning and purpose* of data, i.e., their value for a specific application [2]. Goal-oriented schemes consider the relevance of information, taking into account the shared context of the communicating agents, timing and bandwidth constraints, and the application-level performance metric that needs to be optimized. Research on the subject gained steam after the development of joint source-channel coding [3] and has since been extended to wider *semantic* aspects [4], is mostly focused on goal-oriented compression. Instead of classical reliability metrics, the semantic approach defines a complex, application-dependent distortion function:

Federico Chiariotti (corresponding author, federico.chiariotti@unipd.it) and Leonardo Badia (leonardo.badia@unipd.it) are with the Department of Information Engineering, University of Padova, Padua, Italy. Andrea Munari (andrea.munari@dlr.de) is with the Institute of Communications and Navigation, German Aerospace Center (DLR), Weßling, Germany. Petar Popovski (petarp@es.aau.dk) is with the Department of Electronic Systems, Aalborg University, Aalborg Øst, Denmark. This work was supported in part by the Italian National Recovery and Resilience Plan (NRRP), as part of the RESTART partnership (PE00000001), under the European Union NextGenerationEU Project, and 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 identification number: 16KISK022. A preliminary version of this work has been presented at the IEEE INFOCOM 2025 conference [1].

even if part of a message is lost, distorted, or omitted, the objective is to convey the intended meaning.

On the other hand, a parallel approach has been developed by the Internet of Things (IoT) community, focusing on *medium access* instead of coding. In this case, the relevance of information depends on the error of a remote monitor that estimates the state of a dynamic process through sensor updates. The accuracy of the estimate will tend to degrade over time, unless new updates are received. Age of Information (AoI), which represents the time elapsed since the generation of the last received status report [5], captures this basic relation [6], but it is only a proxy for the actual relevance of sensory information, which depends on the stochastic evolution of the process. The Value of Information (VoI) is a more recent metric that directly considers goal-oriented aspects by measuring the estimation error directly, allowing for more context-aware access schemes, but also increasing their complexity. In order to capture both the need for fresh updates and their relevance [7], the Age of Incorrect Information (AoII) considers a linear penalty counting the time elapsed since the last variation of system conditions [8].

The design of medium access schemes that can minimize AoI or AoII is an important problem in goal-oriented communication, as the relevance of sensor information is known to individual nodes, requiring a distributed approach. This is particularly relevant in scenarios with a large number of sensors and relatively rare events in each location, such as anomaly tracking [9]: scheduled schemes can minimize AoI, or even the expected VoI [10], but the centralized scheduler cannot be aware of anomalies, leading to a higher AoII. However, most of the relevant literature still considers centralized setups due to the need to coordinate transmissions [11] to avoid the collision issue that plagues classical random access protocols such as ALOHA [12], even when using feedback from the common receiver [13] to resolve collisions by computing the state of other contending nodes [14].

The study of random access protocols that can act in a truly goal-oriented fashion, minimizing AoII and fully exploiting the knowledge that centralized schemes lack, is still in its infancy [15], as the analysis of AoII is complex even for simple ALOHA-based protocols [16], [17]. Distributed protocols that can take the content of sensor observations into account are rare in the relevant literature: while a centralized controller cannot exploit the knowledge of the sensors’ true observations, distributed protocols are often plagued by collisions [18]. Sensors can decide whether and when to transmit based on their own observations, but they do not know what other

sensors are observing, and which decisions they might make as a result. This work aims at filling this gap by designing a distributed scheme that uses Dynamic Epistemic Logic (DEL) [19] to allow nodes to employ deductive reasoning over others' states based on *common knowledge* information about their behavior. This can reduce both the frequency of collisions [20] and the time needed to resolve them [21].

We design Dynamic Epistemic Logic for Tracking Anomalies (DELTA), a protocol that adopts DEL to allow sensors to minimize AoII distributedly. Each node considers its belief that it is the one with the highest AoII and then acts accordingly: listening to acknowledgments (ACKs) guarantees that it is able to track everyone else's AoI, using this information to update its belief over others' AoII. The protocol considers a simple binary relevance model, which can however represent a variety of applications, such as (i), a set of wireless sensors reporting anomalies, e.g., excessive temperatures in a factory setting, to a common access point, in which the sensor detecting the occurrence of an anomaly remains in an alert state until it successfully reports it [22], or (ii), a scenario in which agents request access to computing resources over a shared channel, sending a request/interrupt to the common computing engine [23] when they receive a task [24].

To the best of our knowledge, we are the first to combine DEL and goal-oriented communication, designing a random access protocol that exploits this information to provide superior performance over scheduled approaches. The contributions of this paper are listed as follows.

- We introduce DELTA, a random access protocol based on inference reasoning, formally proving that it can allow multiple sensors to efficiently operate in a goal-oriented fashion based on common knowledge information;
- We analyze the protocol settings, providing an exact optimization framework for the collision resolution phase of the protocol and an approximate semi-Markov model for the epistemic reasoning phase;
- We provide an analysis of the effects of various feedback models, showing that the protocol is robust to errors in the feedback channel, degrading gracefully even in very difficult scenarios.

DELTA can reduce the probability that the AoII is over a set threshold by 30–80% with respect to scheduled schemes if the offered load is below 0.5, achieving much better performance than existing random access schemes, as well as reducing the expected AoII with respect to both scheduled and random access alternatives. A preliminary version of this work was presented as a conference paper [1]. There are two major contributions in this work compared to [1]. First, we design a collision resolution scheme that is more advanced than the one in [1], with a superior performance under ideal feedback. Second, we analyze the impact of imperfect feedback. Several feedback models are introduced for this purpose. The results confirm the robustness of DELTA with respect to different imperfect feedback scenarios.

The rest of this paper is organized as follows: first, Sec. II presents the state of the art. Sec. III then defines the communication system model, and the DELTA protocol is specified in Sec. IV, along with the theoretical analysis of its parameters.

We then describe the simulation results in Sec. V, while Sec. VI concludes the paper and presents some possible avenues of future work.

II. RELATED WORK

The analysis of AoII and other AoI extensions in distributed settings is still in its infancy. The existing random access schemes that target information freshness, either require a certain side coordination, or a traffic is extremely sporadic [18], [25]. Even though it was studied in the seminal paper that first defined AoI [5], where the metric was originally introduced for vehicular networks, relatively few works have explicitly considered medium access. A common approach is to treat centralized coordinated access [10], [26], due to the complexity of keeping track of the system state in distributed schemes, as well as information locality: since sensors operate without knowing what the others measure, the collision risk becomes acute unless access is centrally scheduled [18]. Several recent studies [27] considering AoI in random access channels point out how collisions have a detrimental effect on AoI, even when considering carrier sensing [20] and collision resolution mechanisms [21]. The efforts to prevent nodes from entering collisions are mostly circumscribed to the threshold ALOHA approach [16], which can be adapted dynamically to time-varying traffic conditions [28]. However, threshold-based methods can be efficient for AoI but are suboptimal for anomaly reporting due to the overhead incurred due to waiting until an AoII threshold is reached [29].

Deterministic access quickly becomes AoI-optimal for large networks [30]; however, this only holds if the traffic is intense. There are very few investigations on the freshness of anomaly reporting, which is not expected to be persistent. Most anomaly tracking applications, where staleness is better quantified by AoII, do not require constant updates and avoid unnecessary transmissions, improving battery lifetime and congestion [10]. Scenarios include vehicular flow management in which critical reporting by a vehicle is not constant and depends on its position [31], environmental supervision in smart agriculture, wildlife tracking, or monitoring for safety and security purposes in domotic, industrial, or smart grid scenarios [32]. Even medical supervision of elderly or chronic patients likely only reports relevant condition changes [33]. In all these scenarios the traffic is intermittent, but far from sporadic (e.g., vehicular communications may require an exchange of data with an update every second or so [34]), and the tracked anomalies are sudden and variable across the users. In this context, analyzing AoII in more complex reservation-based protocols is often only possible as the number of nodes grows to infinity [35], while precise results for finite networks have been provided just for simple schemes, such as ALOHA [36]. To the best of our knowledge, the only work to actively optimize AoII instead of analyzing existing schemes is [15], whose results are still inferior to simple round-robin.

We then consider the work on epistemic logic, a branch of formal reasoning dealing with the inference, transfer, and update of knowledge among multiple agents [37], [38]. When knowledge evolves over time and successive interactions, this

is referred to as DEL, and finds applications in social networks and cryptography [39]. The solution is often obtained through meta-reasoning on whether *other* agents are able to solve the problem. For example, in the well-known “muddy children puzzle,” agents may possess an individual trait (i.e., a dirty face) or not. This information is not directly available, as each agent only knows if others have the trait, and that at least one child does [40]. Proceeding by induction, one can determine the exact number of muddy faces over a few rounds.

There have been a few attempts at introducing DEL at the network level, mostly driven by the use of AI-empowered devices. For example, [32] discusses the ability of IoT systems to combine local knowledge of individual nodes through automated reasoning, so as to gain further meta-information. Quite recently, [38] has explored AI for network virtualization, and leverages epistemic logic to improve over the uncertainties of AI with respect to traditional software-based virtual network functions. However, none of these or other similar proposals consider DEL for medium access.

III. SYSTEM MODEL

Consider a discrete-time system with a set \mathcal{N} of sensors (also referred to as nodes), each of which measures an independent quantity and can detect anomalies. We denote the number of nodes as $N = |\mathcal{N}|$ and the state at time step t as $\mathbf{x}_t \in \{0, 1\}^N$, whose n -th component $x_{n,t}$ corresponds to the state of sensor n at time t . The free evolution of the process monitored by sensor n is driven by λ_n , the transition probability from the normal state 0 to the anomalous state 1. The sensor starts in the normal state, and once an anomaly occurs, it remains in state 1 unless an external command resolves the anomaly. The transition matrix \mathbf{A}_n of the absorbing Markov chain is then

$$\mathbf{A}_n = \begin{pmatrix} 1 - \lambda_n & \lambda_n \\ 0 & 1 \end{pmatrix}. \quad (1)$$

However, we assume that the gateway can receive warnings from the sensors and adjust its estimate of the system state or send commands to actuators to resolve the anomaly. This effectively means that, upon a successful transmission, a sensor returns to the normal state. We define the indicator variable $s_{n,t} \in \{0, 1\}$, which is equal to 1 if n successfully transmits at time t and 0 otherwise. We then define the time-dependent transition matrix $\mathbf{B}_{n,t}$:

$$\mathbf{B}_{n,t} = \begin{pmatrix} 1 - \lambda_n(1 - s_{n,t}) & \lambda_n(1 - s_{n,t}) \\ s_{n,t} & 1 - s_{n,t} \end{pmatrix}. \quad (2)$$

In other words, while \mathbf{A}_n represents the natural transition of the process, which is absorbing, $\mathbf{B}_{n,t}$ includes the reporting process as part of the evolution of the system. Whenever the sensor is in state 1 or an anomaly occurs, it can go back to state 0 by informing the gateway.¹ We then define the AoI of node n at time t , denoted as $\Delta_{n,t}$, as

$$\Delta_{n,t} = t - \max_{\tau \in \{1, \dots, t\}} \tau s_{n,t-\tau}. \quad (3)$$

¹For the sake of simplicity, we consider transmissions to be instantaneous. The case in which transmissions incur a delay of 1 slot can be dealt with by adding 1 to all AoI and AoII measurements in the following.




Fig. 1. Example of the AoI and AoII evolution for a node.

However, AoI is not meaningful in our case, as a sensor might spend a long time with nothing to report: as long as its state is normal, new updates from it are not necessary. We then introduce the AoII $\Theta_{n,t}$ [8], which is defined as

$$\Theta_{n,t} = t - \arg \max_{\theta \in \{t - \Delta_{n,t} + 1, \dots, t\}} \theta x_{n,t-\theta}. \quad (4)$$

As Fig. 1 shows, the AoI grows even while in the normal state, while the AoII only grows in the anomalous state. We also define the AoII violation probability $V_n(\theta^*)$ as

$$V_n(\theta^*) = \mathbb{E} \left[\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^T \mathcal{I}(\Theta_{n,t} > \theta^*) \right]. \quad (5)$$

where $\mathcal{I}(\cdot)$ is the indicator function, equal to 1 if the condition in the argument is true and 0 otherwise.

We consider the wireless communication system to operate in Time Division Duplex (TDD) mode, so that each time slot is divided in an uplink and downlink part. During the uplink part, each sensor may transmit or remain silent. The uplink is modeled as a collision channel, in which transmissions are never successful if more than one node is active. If a single node n transmits, its packet erasure probability is ε_n .

During the downlink part, all sensors are in listening mode. If the uplink transmission was successful, the ACK packet from the gateway informs all nodes of the identity of the transmitter, while if it was unsuccessful, either because of a collision or a wireless channel erasure, a Negative ACK (NACK) packet informs all nodes of the failure, but does not report the identity of the transmitting nodes. Finally, if no node transmitted, the gateway is silent [41].

We will consider four different models for the ACK and NACK transmission channel from the gateway to the nodes:

- An *ideal* feedback channel, in which all nodes receive the messages without errors;
- A *noisy* feedback channel, in which ACKs and NACKs are always distinguished, but the decoded identity of the intended recipient of the ACK is a Gaussian random variable with a standard deviation σ_f , as explained below;
- An *erasure* feedback channel, in which each node may be unable to decode the ACK with probability ε_f , but knows whether a feedback message was sent;
- A *deletion* feedback channel, in which a node is unable to even know if a feedback message was transmitted or not with probability ω_f .

In general, the protocol is robust to an imperfect feedback channel, and we will discuss the countermeasures to deal with this case in the following. The noisy model is inspired by new IoT technologies such as wake-up radio: recent studies

and standardization efforts by 3GPP [42] show that a gateway using simple on-off keying can deliver a 16 bit ACK in 0.1 ms [43], and the use of more advanced modulations like Orthogonal Frequency Division Modulation (OFDM) can further improve their efficiency [44]. Accordingly, the duration of the downlink phase can be safely neglected, and the electronics implementing the receiver can be designed to consume orders of magnitude less than a standard radio. In this case, confusing ACKs and NACKs becomes almost impossible, as the code can be designed for a wide separation of the two, but the duration of the ACK signal may be misinterpreted by nodes, leading to a certain probability of error over the node ID. In this case, we consider a Gaussian noise over the decoded ID, $w \sim \mathcal{N}(0, \sigma_f^2)$: if node n receives an ACK for a packet sent by node m , the decoded ID is

$$\hat{m}_n = \text{mod}(\text{int}(m - 1 + w), n) + 1, \quad (6)$$

where $\text{mod}(m, n)$ is the integer modulo function.

Finally, if node n transmitted during the slot, it will always assume that an ACK is meant for its own packet independently of the noise, as only one packet can be acknowledged in a given slot. On the other hand, the erasure and deletion models correspond to more classical digital feedback channel models, in which the nodes are in receive mode during the downlink phase of each round. This usually ensures a very low feedback error probability, as the gateway can transmit using a high power and a robust modulation and coding, but requires a higher energy expense for the nodes.

We denote the feedback (or lack of it) received by node n after slot t as $f_{n,t}$, and the history of feedback up to time t as $\mathbf{f}_{n,0:t} \in \mathcal{F}_t$. Our objective is then to find a transmission strategy $\pi_{n,t}^* : \mathcal{F}_t \times \mathbb{N} \rightarrow [0, 1]$ matching the past history of received feedback and the current AoII to a transmission probability and achieving

$$\pi_{n,t}^* = \arg \min_{\pi: \mathcal{F}_t \times \mathbb{N}} V_n(\theta^*). \quad (7)$$

As the space of possible strategies is extremely large, our solution relies on heuristic reasoning.

In the following, we will refer to random variables using capital letters, e.g., X , while their realizations will use the corresponding lowercase letter, e.g., x . The Probability Mass Function (PMF) of X will be indicated as $p_X(x)$, and the corresponding Cumulative Distribution Function (CDF) will be $P_X(x)$. Vectors are indicated as bold lowercase letters, e.g., \mathbf{x} , whose n -th element is denoted by x_n . Matrix symbols are bold capital letters, e.g., \mathbf{A} , whose element in row m and column n is denoted by $A_{m,n}$.

IV. THE DELTA PROTOCOL

In order to reduce the AoII violation probability, we propose the Dynamic Epistemic Logic for Tracking Anomalies (DELTA) protocol, which is based on the notion of *common knowledge* as defined in [19]. DEL is a formal framework to describe the dynamics of beliefs in multi-agent systems, which distinguishes between general and common knowledge proposition. A proposition is general knowledge if its truth value is known to all agents, while for it to be common




Fig. 2. DELTA state diagram.

knowledge, the fact that it is general knowledge also needs to be known to all agents, extending recursively to infinity. The use of common knowledge-based Bayesian reasoning allows DELTA nodes to maintain a shared understanding of the state of the system, which each sensor can combine with its own private observations to make communication decisions.² Furthermore, the public outcome of these decisions can be used by sensors to infer other nodes' private knowledge, following a Bayesian framework. The crucial aspect to enable this is the public nature of ACKs. In the following, we only consider the ideal and noisy feedback channel cases, but we adapt DELTA to an imperfect feedback channel in Sec. IV-E.

A. Protocol Definition

The DELTA protocol includes 4 phases, and transitions between them only depend on publicly available information, e.g., the outcome of the previous slot. The DELTA phase diagram is shown in Fig. 2.

The *Zero-Wait (ZW)* phase is the default mode of operation: during this phase, each sensor transmits whenever its state changes, i.e., an anomaly occurs. This allows us to keep the AoII equal to 0 when the system is empty. Sensors remain in this phase until a transmission fails due to multiple sensors simultaneously observing anomalies or a wireless channel erasure. As the gateway transmits a NACKs signal to inform sensors of the collision, all sensors switch to the *Collision Resolution (CR)* phase [21], recording their membership in the collision set through an indicator variable $m_{n,t}$.

During the CR phase, nodes with $m_{n,t} = 0$ never transmit. In the first slot after the collision, members of the collision set transmit with a certain probability p . In the following slots, the nodes keep transmitting with the same probability until there is a successful transmission, i.e., an ACK is received: in this case, the nodes transition to the *Collision Exit (CE)* phase. During this phase, nodes that are not in the collision set remain silent, while the node that successfully transmitted exits the collision set by setting $m_{n,t} = 0$. All remaining members of the collision set transmit with probability 1. This strategy does not always lead to an immediate reduction of the AoII. While it is very effective if there are 1 or 2 nodes in the collision set, it always results in a collision if there are more than 2. However, it also has a significant advantage over less aggressive strategies: whatever the outcome of the

²In the following, we assume that vector λ is known to all nodes. If this is not the case, empirically estimating the activation rate of other sensors without signaling is easy if the feedback includes the AoII of the successfully reported anomaly, as the distribution of inter-anomaly times is geometric.

slot, it allows all nodes to know the phase of the system. If the CE phase is successful, i.e., the slot is either silent or a correct transmission, the feedback implicitly tells the nodes that the collision set is empty, and they should switch to the *Belief Threshold (BT)* phase. On the other hand, the second collision allows all nodes to know that the initial collision is still unresolved, and that there should be another CR phase. This allows for the preservation of common knowledge, which can provide significant long-term benefits.

Finally, the BT phase allows sensors to gradually reduce the maximum estimated AoII, eventually going back to the ZW phase: as the sequence of CR and CE phases can take several steps, anomalies may have accumulated, and several sensors may have a high AoII. Consequently, the sensors need to get back to a state in which they have common knowledge that everyone is in state 0 before ZW operation can safely resume.

Let us denote the highest possible AoII that a node might have given the common knowledge information as $\psi_{n,t}$:

$$\psi_{n,t} = \sup \{ \theta \in \mathbb{N} : P(\mathbf{f}_{n,0:t} | \Theta_{n,t} = \theta) \neq 0 \}. \quad (8)$$

Under an ideal feedback model, all nodes can compute $\psi_{n,t}$ and obtain the same result, and they know that $\Theta_{n,t} \leq \psi_{n,t} \quad \forall t, n$. The computational complexity of this calculation is $O(\psi_{n,t})$, so the overall computational complexity of the maximum AoII update is $O(N\psi_{\max})$, where $\psi_{\max} = \max \psi_t$. Node n 's AoII $\Theta_{n,t}$ is the highest if no node has higher AoII, and the activation of each node is independent. The probability that node n has the highest AoII, given the vector ψ_t , is then

$$z_{n,t}(\Theta_{n,t}, \psi_t) = \prod_{m \neq n} (1 - \lambda_m)^{[\psi_{m,t} - \Theta_{n,t} + 1]^+}, \quad (9)$$

where $[x]^+ = \max(0, x)$ is the positive part operator. In the BT phase, we set a threshold Z , and node n transmits with probability 1 if $z_{n,t} > Z$. If $\psi_t = \mathbf{0}_N$, i.e., the all-zero vector of length N , the system goes back to the ZW phase. The complexity of the calculation of $z_{n,t}(\Theta_{n,t}, \psi_t)$ is $O(N\psi_{\max})$. As all other DELTA operations are $O(1)$, the overall complexity of the protocol is $O(N\psi_{\max})$, which is manageable even for rather large networks.

We note that collisions are more common in the BT phase than in ZW, as nodes must be more aggressive to gradually reduce ψ_t . All collisions are handled identically, regardless of the phase during which they originated. The full decision-making algorithm for each sensor is presented as Alg. 1.

Finally, we provide the following theorem about the correctness of the protocol:

Theorem 1. *Under an ideal feedback model, the phase and the value of ψ_t are always common knowledge, and the phase is always the same for all nodes.*

The proof of the theorem, involving a DEL model of the network, is presented in the Appendix.

B. Collision Resolution Phase Optimization

The expected number of slots τ_c required to resolve a collision depends on the number C of colliding nodes, which transmit with the same probability p until the collision is

Algorithm 1 Pseudocode of the DELTA protocol

```

Require: phase,  $Z$ ,  $\mathbf{p}$ ,  $x_{n,t}$ , NACK,  $m_{n,t}$ ,  $c_t$ ,  $\psi_{t-1}$ 
1: if NACK then
2:   if phase = CE then
3:      $c_t \leftarrow c_t + 1$ 
4:   phase  $\leftarrow$  CR
5: if ACK and phase = CR then
6:   phase  $\leftarrow$  CE
7: if phase = BT then
8:    $\psi_t \leftarrow$  UPDATEMAXIMUMPOSSIBLEAOII( $\psi_{t-1}$ )
9:   if  $\max(\psi_t) = 0$  then
10:    phase  $\leftarrow$  ZW
11: if phase = CE and (not NACK) then
12:   phase  $\leftarrow$  BT,  $c_t \leftarrow 0$ 
13: if  $x_{n,t} = 0$  then
14:   return 0
15: else
16:   switch phase do
17:     case ZW: return 1
18:     case CR: return  $m_{n,t} \times$  BERNOULLISAMPLE( $p(c_t)$ )
19:     case CE: return  $m_{n,t}$ 
20:     case BT: return HIGHESTAOIIPROB( $\theta_t, \psi_t$ )  $> Z$ 

```

resolved. As defined in Sec. III, the wireless channel erasure probability for node n is denoted as ε_n . The probability of success in any given slot when there are c colliding nodes is

$$\sigma(c, p, \varepsilon_n) = (1 - \varepsilon_n) \text{Bin}(1; c, p), \quad (10)$$

where $\text{Bin}(k; N, p) = \binom{N}{k} p^k (1-p)^{N-k}$ is the binomial PMF. After the first ACK, the remaining colliding nodes transmit with probability 1 in the CE phase. This means that $C - 1$ nodes will collide if $C > 2$. We then define vector \mathbf{p} , whose i -th element represents the transmission probability in the i -th collision resolution phase.

If all nodes have the same ε , we can represent the cycle starting from c colliding nodes as an absorbing Markov chain with c states, representing each individual CR phase. The transition from one state to the next is the CE phase, and the structure of the protocol prevents the size of the collision set from increasing. The transition probability matrix is

$$\mathbf{P}_c = \begin{pmatrix} \mathbf{B} & \sigma(2, p_{c-1}, \varepsilon) \mathbf{u}_{c-1}^{c-1} \\ (\mathbf{0}_{c-1})^\top & 1 \end{pmatrix}, \quad (11)$$

where $\mathbf{0}_N$ is the column vector of length N whose elements are all equal to 0, and the non-zero elements of matrix \mathbf{B} are

$$B_{ij} = \begin{cases} 1 - \sigma(c - i + 1, p_i, \varepsilon), & j = i; \\ \sigma(c - i + 1, p_i, \varepsilon), & j = i + 1. \end{cases} \quad (12)$$

The time τ_c until absorption, i.e., until the collision is fully resolved, follows a discrete phase-type distribution characterized by the matrix \mathbf{P}_c . The CDF of τ_c is simply given by the corresponding element of the t -step matrix, $p_{\tau_c}(t) = (\mathbf{P}_c)_{1,c}^t$. In the case where $c = 1$, i.e., when a single node's transmission failed because of the channel, the time until absorption reduces to a geometric random variable, i.e., $\tau_1 \sim \text{Geo}(p_1)$.

Theorem 2. *If the colliding set was a singleton, i.e., $C = 1$, the expected duration of the subsequent CR and CE cycle is*

$$\mathbb{E}[\tau_1] = 1 + ((1 - \varepsilon)p_1)^{-1}. \quad (13)$$

For a set of $c > 1$ colliding nodes with the same ε , the expected duration of a cycle of CR-CE phases, which begins after the initial collision and ends when the collision set is empty, is

$$\mathbb{E}[\tau_c] = c - 1 + \varepsilon + \frac{\varepsilon}{(1 - \varepsilon)p_c} + \sum_{i=0}^{c-2} \frac{1}{\sigma(c - i, p_{i+1}, \varepsilon)}. \quad (14)$$

Proof: We begin by proving the theorem in the singleton case, in which there is a single CR phase, whose duration is geometrically distributed with parameter $(1 - \varepsilon)p_1$. An additional slot needs to be added to account for the CE phase.

In the general case, the expected time until absorption of a Markov chain is hard to compute, but the structure of the transition matrix simplifies the problem. Any state i is reached from $i - 1$ with a successful transmission after a geometrically distributed number of failures, i.e., self-transitions:

$$\mathbb{E}[\tau_{i-1,i}|C = c, p_{i-1}] = (\sigma(c - i, p_{i-1}, \varepsilon))^{-1}. \quad (15)$$

The number of self-transitions in each state is independent from what happens in other states due to the Markov property, and the protocol requires $c - 1$ CR phases to reach the absorbing state c . Additionally, there are $c - 2$ collisions caused by the intermediate CE phases, during which the nodes discover that the collision set is not empty. Finally, there is one more CE phase from the last colliding node after reaching state c . If the transmission is successful, the cycle is over, but if there is a wireless channel loss, there is one more singleton collision resolution cycle after it. ■

However, the value of C is unknown to the sensors. If we consider the ZW phase in a system in which all sensors have the same activation probability λ , we get

$$p_C(c|ZW) = \text{Bin}(c; N, \lambda) [1 - (1 - \varepsilon)\delta(c, 1)], \quad (16)$$

where $\delta(m, n)$ is the Kronecker delta function, equal to 1 if the two arguments are equal and 0 otherwise. We can also easily get the total failure probability $p_f(ZW) = \sum_{c=1}^N p_C(c|ZW)$. We can then apply the law of total probability, adding the $c - 1$ CE phases as in Theorem 2, to obtain the CDF of the duration of a collision resolution cycle, with $\eta_c = 1 - (1 - \varepsilon)p_c$:

$$P_\tau(t|ZW) = \frac{\varepsilon(1-\varepsilon)}{p_f(ZW)} \left[\text{Bin}(1; N, \lambda) (1 - \eta_1^{t-1}) \right. \\ \left. + \sum_{c=2}^{\min(N,t)} \text{Bin}(c; N, \lambda) \left(\frac{(\mathbf{P}_c)_{1,c}^{t-c+1}}{\varepsilon} + \sum_{k=1}^{t-2c+1} (\mathbf{P}_c)_{1,c}^{t-c-k} \eta_c^{k-1} p_c \right) \right]. \quad (17)$$

Theorem 3. The optimization problem defined by

$$\mathbf{p}^* = \arg \min_{\mathbf{p} \in [0,1]^N} \sum_{c=1}^N p_C(c|ZW) \mathbb{E}[\tau_c], \quad (18)$$

is convex if all nodes have the same λ and ε , and the optimal transmission probability p_i^* for the i -th round of CR is the unique solution of

$$\frac{\text{Bin}(1; N_i, \lambda)\varepsilon}{(p_i^*)^2} + \sum_{c=2}^{N_i} \text{Bin}(c; N_i, \lambda) \frac{1 - cp_i^*}{c(p_i^*)^2(1 - p_i^*)^c} = 0, \quad (19)$$

where $N_i = N - i + 1$. In the N -th CR phase, $p_N^* = 1$.

Proof: Since each CR phase is independent from all others, each element of \mathbf{p} can be separately optimized to minimize the expected duration of that individual phase. We then take the probability for the first CR phase as a representative example, obtaining problem

$$p_1^* = \arg \min_{p \in [0,1]} \left[\sum_{c=1}^N w_c (\sigma(c, p, \varepsilon))^{-1} \right], \quad (20)$$

where $w_c = p_C(c|ZW)$, which is independent from p_1 . In order to prove that it is convex, we only need to prove that each component of the weighted sum is convex. The first one, with $c = 1$, is proportional to p^{-1} , so it is convex for $p > 0$. We show that components with $c > 1$ are convex by taking the second derivative of $(\sigma(c, p, \varepsilon))^{-1}$ with respect to p :

$$\frac{\partial^2 (\sigma(c, p, \varepsilon))^{-1}}{\partial p^2} = \frac{c(c+1)p^2 - 2(c+1)p + 2}{(1 - \varepsilon)cp^3(1 - p)^{c+1}}. \quad (21)$$

As $c > 1$, $p \in (0, 1)$, and $1 - \varepsilon$ is always positive, so is the denominator. The second derivative is then positive if

$$c(c+1)p^2 - 2(c+1)p + 2 > 0. \quad (22)$$

This quadratic equation has no real solution for $c > 1$. The objective of the problem in (18) is then convex, and the set $[0, 1]^N$ is also convex, proving the first half of the theorem.

We can then trivially show that (19) is the first derivative of the objective in (20). We can trivially prove that the two extremes, $p = 0$ and $p = 1$, lead to an infinite expected duration for $N > 1$: if $p = 0$, no node ever transmits, while if $p = 1$, the nodes will keep colliding forever whenever the remaining collision set is not a singleton [12]. The maximum is then inside the interval for $N > 1$.

Finally, we can prove that (19) is a multiple of the first derivative of the optimization function in (18), and finding its root in $(0, 1)$ is equivalent to finding the minimum. ■

Since the problem is convex, there exists a single global optimum in the optimization set, which can be found by setting the derivative to 0, i.e., by the condition in (18). However, solving the equation is not possible in closed form, as it involves a hypergeometric function. However, a solution can be easily found by applying the bisection method [45]. This can be then stored efficiently as a look-up table to avoid a further computational burden.

C. DELTA+

A fixed transmission probability still does not fully account for the information received through public announcements: each failed or silent slot can be used as a Bayesian update. This principle was adopted as part of the Sift protocol [46], which provided an optimal solution for a known number of colliders and an approximated one with an unknown number. In our case, the initial distribution of the number of colliders in the first CR phase is

$$\phi_0(c) = \mathbb{1}(c=1) \frac{(1 - (1 - \varepsilon)\delta(c, 1)) \text{Bin}(1; N, \lambda)}{\varepsilon \text{Bin}(1; N, \lambda) + \sum_{c'=2}^N \text{Bin}(c'; N, \lambda)}, \quad (23)$$




Fig. 3. Approximated semi-Markov model of DELTA with $K=3N$, $\Psi=4$.

where $\mathbb{1}(x)$ is the stepwise function, equal to 1 if $x \geq 0$ and 0 otherwise. We can then update the belief distribution after an ACK by applying Bayes' theorem:

$$\phi_{j+1}^{\text{CR}}(c|\text{ACK}) = \frac{\phi_j(c+1) [(c+1)p_j(1-\varepsilon)(1-p_j)^c]}{\sum_{c'=0}^N \phi_j(c') [c'p_j(1-\varepsilon)(1-p_j)^{c'-1}]} \quad (24)$$

After a silent slot, we get

$$\phi_{j+1}^{\text{CR}}(c|\text{SIL}) = \frac{\phi_j(c)(1-p_j)^c}{\sum_{c'=0}^N \phi_j(c')(1-p_j)^{c'}} \quad (25)$$

Finally, we can perform a similar update after a NACK:

$$\phi_{j+1}^{\text{CR}}(c|\text{NACK}) = \frac{\phi_j(c)p_{\text{NACK}}(c)}{\sum_{c'=0}^N \phi_j(c')p_{\text{NACK}}(c')} \quad (26)$$

where $p_{\text{NACK}}(c)$ is

$$p_{\text{NACK}}(c) = 1 - (1-p_j)^c - cp_j(1-\varepsilon)(1-p_j)^{c-1} \quad (27)$$

After an unsuccessful CE phase, we update the belief as

$$\phi_{j+1}^{\text{CE}}(c|\text{NACK}) = \frac{\phi_j(c)\mathbb{1}(c-1)(1-(1-\varepsilon)\delta(c,1))}{\varepsilon\phi_j(1) + \sum_{c'=2}^N \phi_j(c')} \quad (28)$$

Using this belief distribution, the optimal transmission probability p_j^* is the solution of

$$\frac{\phi_j(1)}{(p_j^*)^2} + \sum_{c=2}^N \frac{(1-cp_j^*)\phi_j(c)}{c(p_j^*)^2(1-p_j^*)^c} = 0 \quad (29)$$

The proof that this solution is optimal trivially follows from Theorem 3. We will refer to the version of the protocol using this slot-level belief update as DELTA+, to distinguish it from the basic version.

In the basic version of DELTA, the transmission probabilities in the CR phase can be pre-computed and stored as a look-up table, with no computational cost at runtime, but DELTA+ must solve (29) at each CR step. If we consider a precision \tilde{p} , the bisection method can solve this equation in $O(\log_2(\tilde{p}^{-1}))$ time. The overall computational complexity of each phase of DELTA and DELTA+ is summarized in Table I.

D. Belief Threshold Optimization

We can create a semi-Markov model of the system, as shown in Fig. 3, by applying some simplifications: firstly, we consider nodes with the same activation probability λ . Setting a threshold Z on the probability of being the highest node then corresponds to setting a maximum number $K = \frac{\log(Z)}{\log(1-\lambda)}$ of possible slots in which the nodes transmit. Secondly, we consider some approximations in the outcomes of the BT phase, which we will discuss below.

The ZW state always leads to a collision, i.e., to a CR phase, but the state of the model also keeps track of the highest ψ^* (which is always 0 for the ZW phase). Correspondingly, each sequence of CR and CE phases ends with a transition to the BT phase, but ψ depends on the duration of the sequence, which we have analyzed above. During the BT phase, we simplify the model by considering the case in which a single collision resolution phase led to the current state, i.e., by discarding secondary collisions that happen while in the BT phase. Given the maximum possible AoII ψ , we can obtain the conditioned PMF of the number of colliders by applying Bayes' theorem:

$$p_C(C|\psi) = p_C(c|\text{ZW})p_{\tau_c}(\psi^*)(p_{\tau}(\psi))^{-1}, \quad (30)$$

where $p_{\tau}(\psi)$ is the PMF corresponding to the CDF in (17).

We then consider a pessimistic and an optimistic model. The pessimistic model considers $L(\psi) = N$, i.e., all nodes are considered as possible colliders, independently of their $\psi_{n,t}$. This is a pessimistic estimate, as some nodes might have a lower $\psi_{n,t}$ such that it is common knowledge that they cannot be part of the collision set. On the other hand, the optimistic model subtracts the expected number of colliders from the set of active nodes, considering that they have a much lower AoI and, as such, will not transmit. This model is optimistic, as it considers a single collision resolution phase, while the previous dynamics might be more complex and lead to a larger number of potential colliders. The number of active nodes in the optimistic model is $L(\psi) = N - \mathbb{E}[C|\psi]$. Each sensor transmits with probability $\beta = 1 - (1-\lambda)^{\frac{K}{L(\psi)}}$, so the collision probability is

$$\xi(\psi) = 1 - (1-\lambda)^K - (1-\varepsilon) \text{Bin}(1; L(\psi), \beta) \quad (31)$$

In the ZW phase, we have $K = 1$. In the BT phase, we typically have less than N active nodes, but we need to set $K > N$, as $\psi_{n,t}$ decreases by $\left\lfloor \frac{K}{L(\psi)} \right\rfloor - 1$ for each BT step, including those whose outcome is a collision. We can also adjust the transmission probability vector \mathbf{p} of a CR cycle following a collision in a BT slot, using $1 - (1-\lambda)^{\frac{K}{L(\psi)}}$ as an activation probability and finding the solution from Theorem 3.

In order to maintain a finite state space \mathcal{S} , we need to set a maximum AoII Ψ , so that $|\mathcal{S}| = 2\Psi + 1$. We can reduce the approximation error as much as possible by considering

TABLE I
COMPUTATIONAL COMPLEXITY OF THE TWO DELTA VARIANTS.

Version	ψ_i update	ZW	CR	CE	BT
DELTA	$O(N\psi_{\max})$	$O(1)$	$O(1)$	$O(1)$	$O(N\psi_{\max})$
DELTA+	$O(N\psi_{\max})$	$O(1)$	$O(N - \log_2(\tilde{p}))$	$O(1)$	$O(N\psi_{\max})$

a large value that will almost never be reached in practice. This analysis can also be used to ascertain the stability of the system: if the steady-state probability of state $CR(\Psi)$ does not decrease as Ψ increases, the system is unstable. We can then give the elements of the transition matrix M of our model, considering the transitions toward state ZW :

$$M_{s,ZW} = (1 - \xi(\psi))\delta(s, BT(\psi))\mathbb{1}(K - \psi L(\psi)). \quad (32)$$

As ψ is reduced by $\lfloor KL(\psi) \rfloor - 1$ steps whenever a collision is avoided in the BT phase, only BT states with a low value of ψ return directly to ZW . We can compute the transition probabilities to CR states as

$$M_{s,CR(\psi)} = \begin{cases} 1, & s = ZW, \psi = 0; \\ \xi(\psi'), & s = BT(\psi'), \psi' = \left[\psi + 1 - \frac{K}{L(\psi')}\right]^+. \end{cases} \quad (33)$$

Finally, we compute the probability of transitioning to the BT phase, considering that ψ is limited to Ψ :

$$M_{s,BT(\psi)} = \begin{cases} \zeta_{\psi'}(\psi - \psi'), & s = CR(\psi'); \\ 1 - \xi(\psi'), & s = BT\left(\psi + 1 - \frac{K}{L(\psi')}\right); \\ \sum_{\ell=\Psi-\psi'}^{\infty} \zeta_{\psi'}(\ell), & s = CR(\psi'), \psi = \Psi; \end{cases} \quad (34)$$

where $\zeta_{\psi'}(\ell)$ is the PMF corresponding to the CDF given in (17), computed using the optimal transmission probability vector $\mathbf{p}^*(\psi')$. However, as the system is not a Markov chain, but a discrete-time semi-Markov model, we have $T_{ZW,CR(0)} = \text{Geo}(\xi(0))$, $T(BT(\psi), s') = 1$, and $T(CR(\psi), BT(\psi')) = \psi' - \psi$. We also consider a pessimistic approximation: if the collision resolution process leads to state $BT(\Psi)$, the time in the CR state will be Ψ , which should be set to a higher value than the time that is reasonably required to resolve a collision. We can easily obtain the steady-state probability distribution β as the solution to the equation $\beta(\mathbf{P} - \mathbf{I}) = 0$, normalized so that $\|\beta\|_1 = 1$. This corresponds to the left eigenvector of M with eigenvalue 1. The steady-state distribution μ is obtained by weighting β by the average sojourn times $\mathbb{E}[T(s, s')]$:

$$\mu(s) = \frac{\sum_{s' \in \mathcal{S}} \beta(s) M(s, s') \mathbb{E}[T(s, s')]}{\sum_{s^*, s^{**} \in \mathcal{S}} \beta(s^*) M(s^*, s^{**}) \mathbb{E}[T(s^*, s^{**})]}, \quad \forall s \in \mathcal{S}. \quad (35)$$

We can then use $\mu(ZW)$ as a proxy for our desired performance and find $K^* = \arg \max_{K \in \mathbb{N} \setminus \{0, 1\}} \mu(ZW)$. Alternatively, we can sum the steady-state probabilities of states that do not violate the AoII requirement.

E. Dealing with Imperfect Feedback and Future Extensions

Theorem 1 requires all nodes to be able to perfectly distinguish between ACKs, NACKs, and silent slots. This condition is met by the ideal and noisy feedback models, as the only confusion in the latter is over the identity of the node receiving the ACK. As we will see in the following, this has a negligible effect on performance, unless the number of nodes in the system is very small.

To compute ψ_t and synchronize phase transitions, all nodes need to receive an ACK or NACK after each communication slot. In the ZW , CR , and CE phases, this issue can be mitigated

by adding only 2 bits to ACK and NACK packets, representing the current phase (with 4 possible values). The gateway knows the outcome of each transmission, as it is the intended receiver. It can then compute the current phase and piggyback it on ACK and NACK packets. This synchronizes the protocol for these three phases where knowing the phase completely determines a node's behavior; unless the same node misses multiple feedback packets, the anomaly will be quickly solved, and the protocol will work as intended. Mitigation is more complex in the BT phase: since computing $z_{n,t}(\Theta_{n,t}, \psi_t)$ requires a full knowledge of what happened in the past, nodes may have slightly different beliefs over the possible states of the system, leading to inconsistent decision-making processes. We will consider a scheme that includes $\max(\psi_t)$ in the feedback packets during the BT phase, while sensors simply remain in the same phase if they do not receive an ACK, relying on the next one to synchronize with the others. This heuristic might not be optimal, but we show that it is robust with respect to feedback errors, as adapting the Bayesian reasoning in the proof of Theorem 1 to this case, considering missed feedback packets as a possible cause of the outcome of each slot, is rather complex.

Additionally, the behavior of the DELTA protocol after a feedback message has been missed is as follows:

- In the ZW and BT phases, the node behaves as if the slot was successful until the next feedback message allows it to synchronize the protocol phase. While this choice is optimistic, it leads nodes to avoid reducing their transmission probability unnecessarily;
- In the CR phase, the node behaves as if the slot failed until the next feedback message allows it to synchronize the protocol phase. In the DELTA+ variant, the belief over the number of colliders is not updated;
- In the CE phase, the node assumes there was a collision, waiting for the next ACK, unless the slot was silent, in which case it moves to the BT phase. In DELTA+, the belief over the number of colliders is not updated.

Under the deletion channel feedback model, nodes in the CE phase always move to the BT phase. The rationale for this design choice is that, while the CR and CE phase involve contention for the channel, and thus minimizing the additional traffic ensures a faster recovery, the other phases of the protocol try to avoid collisions at all costs, and thus increasing the traffic slightly by behaving more aggressively for a short time will not have a significant effect. Additionally, even causing a collision will trigger a NACK, leading most nodes to synchronize their protocol phase.

In the design of DELTA, we selected a collision channel model was selected both for its simplicity and because it puts random access protocols at a disadvantage with respect to models with capture or Successive Interference Cancellation (SIC). However, the modifications to deal with imperfect feedback can also easily deal with multi-packet reception channels, in which the desynchronization between nodes' beliefs is due to the capture of collided packets. This scenario would require an adjustment of the CR phase probability optimization, which would need to account for the higher success probability, but




Fig. 4. $\mu(\text{ZW})$ as a function of K .

we expect DELTA to improve its performance in this scenario. However, the analysis is left as future work.

Additionally, the current model considers sensor observations to be accurate, and anomalies to be solvable only when reported; richer anomaly models, which might entail keeping track of beliefs over the system state at the gateway, are also a possible extension of the basic principles behind DELTA.

V. SIMULATION SETTINGS AND RESULTS

This section presents the results of the Monte Carlo simulations meant to validate the performance of the DELTA protocol. Each setting was tested over a simulation lasting 10^6 slots. In the following, the maximum offered system load $\rho = \|\lambda\|_1$ will be considered as the key parameter.³

A. DELTA Optimization and Robustness

First, we analyze the correctness of the theoretical model and the optimization of the DELTA protocol parameters.

Fig. 4 shows the value of $\mu(\text{ZW})$, which we can use as a proxy for the stability of the protocol, as a function of the chosen K . We used a Monte Carlo simulation to verify the two approximations, and considered a case with a 20% offered load and a case with a 50% offered load. In both cases, the two semi-Markov models lead to the correct optimization of K . However, Fig. 4a shows that the optimistic model tends to be less accurate when the load is low. This is due to the nature of collisions in this case: most of the time, higher values of ψ will be reached due to multiple collisions between few nodes or even wireless channel losses, leading the estimated value of $L(\psi)$ to be too low. In this case, the pessimistic model, which assumes that all nodes have the same $\psi_{n,t}$, is closer to the real results. On the other hand, the opposite is true when $\rho = 0.5$, as shown in Fig. 4b: when the offered load is high, multiple collisions may cause large differences in the nodes' $\psi_{n,t}$ values, so that the pessimistic model foresees a very low probability of remaining in the ZW phase. In this case, even the optimistic model is too conservative when K is high, as collisions will be frequent enough that nodes will have very different values of $\psi_{n,t}$, but it manages to capture the trend up to the optimal value of K , and as such, it can provide a good guideline for system optimization. DELTA is stable with respect to both K and p , and thus robust to errors in

³The code for the protocol and the simulations in this paper is available at https://github.com/signetlabdei/delta_medium_access




Fig. 5. Optimal transmission probability for each consecutive CR round for different values of ρ , with $N = 20$ and $\varepsilon = 0.05$.

the estimation of ρ and ε . In the following, we will show the performance of DELTA with optimized parameters, as well as a version with a fixed value $K = \frac{5}{2}N$, to prove that fixed general settings can perform well in a variety of scenarios.

We can also consider the robustness of the parameter choice in the CR phase: Fig. 5 shows the result of the transmission probability optimization for different load values. We can note that, aside from the case with $\rho = 0.15$, the difference between the outcomes is less than 0.05 for all CR rounds: this means that even significant errors in the load estimation will still lead nodes to behave in a very similar way, resulting in a good protocol performance even under parameter uncertainty.

Finally, we examine the overall behavior of DELTA and DELTA+ under the settings we will consider to evaluate their performance. Fig. 6a-b show the CDF of τ_c , i.e., the duration of a CR/CE cycle before the collision is resolved. While DELTA+ adopts a more complex strategy, it only gains a small advantage for $\rho = 0.75$, and the resolution times are nearly identical for smaller values of the load. Although a lower load improves collision resolution performance, the two versions of the protocol are able to resolve more than 90% of collisions within 10 steps at all load levels. Fig. 6c-d show the CDF of the highest maximum AoII. The two protocols are still very similar, but DELTA+ manages to reduce the right tail of the distribution for higher loads, gaining a small advantage. We note that DELTA spends most of the time in the ZW phase when the load is low, while loads over 0.5 are almost never in that phase. This is a key parameter to determine performance, and violation probability in particular, as the ZW phase is the one in which DELTA is most reactive to new anomalies.

B. Benchmark Protocols

We consider two common centralized scheduling algorithms and three distributed protocols as benchmarks to test the DELTA protocol's performance against them in terms of worst-case AoII minimization. Firstly, we consider *Round-Robin (RR)*, the simplest possible scheduling algorithm. It entirely avoids collisions and does not require sensors to listen to feedback packets, as long as they maintain synchronization, but may lead sensors to wait for a long time if the network is large, as the average AoI is $\frac{N}{2}$ even with an error-free channel [30]. Round-Robin (RR) is also vulnerable to wireless channel losses, as a lost packet needs to wait for a full round before being retransmitted. We also implement a *Maximum Age First (MAF)* strategy, which is commonly adopted in the




Fig. 6. Analysis of the behavior of DELTA and DELTA+ (considering $K = 50$) with $N = 20$ and $\varepsilon = 0.05$.

AoI literature, as it can optimize the average age in multi-source systems [26]. In our case, it is equivalent to RR if $\varepsilon = 0$, and has the same issues in large networks with many sensors, but it can efficiently deal with wireless channel losses by retransmitting the lost packet immediately. However, this requires all sensors to listen to feedback packets, as they need to know when packet losses occur.

The three distributed algorithms are a variation on the ZW policy, with different collision resolution mechanisms. Firstly, nodes with information to send under the *Pure Zero-Wait* (ZW) policy immediately do so with a certain probability p_1 . If their packets are lost, either due to the wireless channel or to a collision, they keep transmitting with the same probability until they receive an ACK and return to the normal state. This corresponds to a classical slotted ALOHA system. We also consider a *Local Zero-Wait* (LZW) scheme with two distinct probabilities. Each node transmits with probability p_1 if it has information to send, then switches to probability p_2 after a failure until the packet is successfully transmitted. This corresponds to a local back-off mechanism after collisions with p_2 -persistence. Both ZW and LZW only require sensors to listen to feedback packets after they transmit.

Finally, the *Global Zero-Wait* (GZW) protocol is similar to LZW, but the back-off mechanism is implemented by all nodes. After a transmission failure, all nodes switch from p_1 to p_2 . They then go back to p_1 after a successful transmission, assuming the collision involved either 1 or 2 nodes. This protocol is fairer than LZW, which can lead colliding nodes to have a lower priority than other nodes with a lower AoII, but requires all nodes to listen to the feedback for every slot.

The values of p_1 and p_2 for the distributed benchmarks were optimized for each specific scenario by performing a grid search over a Markov representation of the protocols.

C. Performance Evaluation: Ideal Feedback

We consider the performance of the protocols under the ideal feedback model by measuring the AoII violation probability $V(\Theta_{\max})$, which corresponds to the fraction of time that the nodes spend with an AoII higher than the threshold value Θ_{\max} . We analyzed the performance with $\Theta_{\max} = 0$, which requires nodes to immediately report anomalies, and $\Theta_{\max} = 5$, which allows for a short delay before the gateway is successfully informed of the anomaly. Unless otherwise stated, we consider a system with $N = 20$ nodes, a channel erasure probability $\varepsilon = 0.05$.

Fig. 7 shows the violation probability and average AoII as a function of the offered load $\rho = \|\lambda\|_1$, i.e., the load on the system if all nodes immediately transmit successfully, which is an upper bound on the actual system load. The plot clearly shows that DELTA outperforms the other random access schemes, which tend to approach the same reliability only for very low values of the offered load. On the other hand, both $V(0)$ and $V(5)$ grow approximately linearly with ρ for Maximum Age First (MAF) scheduling: as expected, centralized scheduling mechanisms can outperform any random access scheme for congested networks, but DELTA manages to outperform MAF for $\rho < 0.55$, which is a significant improvement over the ZW benchmark, as well as a very intense traffic for anomaly reporting applications. The performance of the optimistic, pessimistic, and fixed (with $K = 50$) variants remains almost the same, and a small difference can be seen only for very high loads. Additionally, the DELTA+ variant is slightly better, but the more intelligent collision resolution mechanism only has a limited effect on the final performance of the protocol. On the other hand, the other random access protocols have a much higher sensitivity to parameter changes, and the jumps for small changes in ρ are due to the quantization of p_1 and p_2 , for which the grid search optimization considered a 0.01 step, and to the higher instability of these protocols, which may lead to snowball effects. Finally, if we consider the expected AoII performance, shown in Fig. 7c, we note the same pattern: while other random access mechanisms quickly become unstable as ρ increases, DELTA outperforms MAF for offered loads up to approximately $\rho = 0.5$, and its performance degrades relatively gracefully even for higher loads. We note that DELTA+ has an advantage over the simpler version, but only at higher loads, for which MAF outperforms them both.

We can also consider the performance of the schemes as a function of the number of nodes N , considering a scenario with a relatively low load ($\rho = 0.3$) and one with a high load ($\rho = 0.5$). As Fig. 8a-c show, the performance of DELTA in the low load scenario tends to improve as the number of nodes grows when considering $V(0)$, while $V(5)$ is close to 0 for all network sizes. The expected AoII is also very close to 0 in all cases. On the other hand, the performance of scheduled algorithms gradually degrades as the network size grows due to the longer duration between subsequent transmission opportunities for the same node, and $V(0)$ is high even for small network sizes, as it is difficult for MAF to immediately report anomalies. Random access schemes perform better: among these, LZW confirms its advantage over ZW and GZW.




Fig. 7. AoII violation probability as a function of ρ , $N = 20$.




Fig. 8. AoII violation and expected AoII as a function of N .

Interestingly, all ZW variants tend to perform better for very small networks, degrading as N grows between 4 and 16 and then improving from that point on: in general, random access schemes should benefit from larger network sizes, as collisions become easier to resolve, but the Markov representation used to optimize the network may be less accurate for small network sizes. However, DELTA still significantly outperforms all other schemes. As for the varying λ , the DELTA+ variant has a negligible improvement over the basic version of the protocol. This also holds in the high load scenario, shown in Fig. 8d-f. Even in this scenario, which is close to its maximum load, DELTA is remarkably robust, and $V(0)$ improves as the network size grows, although $V(5)$ tends to increase for larger networks. However, DELTA far outstrips other random access protocols, which quickly collapse due to the high load, and has a significant performance advantage over scheduled schemes for $N > 10$. The performance gain of DELTA actually increases as the network size grows, particularly with respect to the expected AoII, shown in Fig. 8f.

Finally, we consider the robustness to errors in the estimated activation rates: we set a load $\rho = 0.5$, and randomly sampled 100 activation probability vectors $\lambda \sim \mathcal{U}\left(\frac{(1-\nu)\rho}{N}, \frac{(1+\nu)\rho}{N}\right)$.

The input to DELTA was then the average vector, while the actual activation rates exhibited growing differences among nodes as ν increased. The resulting AoII violation probability is shown in Fig. 9: all protocols are robust to this type of disruption, and in particular, DELTA and DELTA+ are insensitive to changes in the activation probabilities if the overall load is approximately correct.

D. Performance Evaluation: Imperfect Feedback

We then evaluate the robustness of the schemes to imperfect feedback, considering the low load ($\rho = 0.3$) and high load ($\rho = 0.5$) scenarios with $N = 20$ and $\varepsilon = 0.05$ and following the three imperfect feedback models outlined in Sec. III.

We first start with the noisy feedback model, in which ACKs, NACKs, and silent slots can always be distinguished, but nodes may erroneously interpret the content of messages: Fig. 10 shows the AoII violation probability as a function of the error standard deviation σ_f . As the figure clearly shows, all protocols except MAF are almost unaffected. On the other hand, MAF is strongly affected by this feedback model, as the feedback messages serve as polling requests: if a node mistakenly believes that it has been polled, it will transmit




Fig. 9. Performance with imperfect rate estimation as a function of the activation probability range ν with $\rho = 0.5$, $N = 20$.




Fig. 10. AoII violation as a function of the feedback noise standard deviation σ_f with $N = 20$.

an update, potentially causing a collision. The performance advantage of DELTA and DELTA+ is unaffected by errors on the feedback, even if they are significant (a standard deviation $\sigma_f = 5$ out of a total of 20 nodes). This is reasonable, as errors on the feedback will affect the belief of nodes only slightly (if at all), considering that there are several nodes that have a high maximum AoII. As the algorithm is very robust with respect to the choice of the belief threshold, errors on the identity of the nodes will have a limited effect.

We then consider the erasure feedback model: Fig. 11 shows performance as a function of the erasure probability ε_f . After including the adaptation of feedback messages discussed in Sec. IV-E, the protocol degrades gracefully in the low load scenario shown in Fig. 11a-b, maintaining a significant advantage over all other schemes. On the other hand, the DELTA+ variant degrades much faster: while we showed that DELTA can deal with some desynchronization, as the transmission probability during the CR phase is extremely robust, DELTA+ relies on much more complex calculations, which lead p to have a much wider range. Consequently, having nodes with different information due to ACK erasure may lead to highly suboptimal outcomes, as nodes might be misled into selecting very low values of p that lead to a significant number of repeated silent slots before the first success allows them to gain

more feedback. Additionally, if a node outside the collision set misses the initial NACK and mistakenly transmits during the CR phase, the resulting collision will lead node in the collision set to decrease their transmission probability, believing the set to be much larger, further compounding the issue. Even if we consider the high load scenario, which is already close to DELTA's saturation point, with collisions becoming a frequent occurrence, the protocol still comes out on top for $\varepsilon_f \leq 0.1$, as shown in Fig. 11c-d. On the other hand, the performance of DELTA+ quickly degrades, becoming even worse than other random access schemes.

Finally, Fig. 12 shows the performance of all schemes under a feedback deletion model: in this case, we only show the scenario with $\rho = 0.5$, as performance is almost identical to the feedback erasure case. The only noticeable difference is that DELTA+ degrades even faster, while the difference with the erasure model is negligible for all other schemes.

VI. CONCLUSION AND FUTURE WORK

In this work, we presented DELTA, a protocol that allows distributed sensor nodes to report anomalies efficiently by relying on the DEL principle of common knowledge information. The protocol considerably outperforms both random access and scheduled schemes under reasonable operating conditions, and its operation is robust to relatively large shifts in its




Fig. 11. AoII violation as a function of the feedback erasure probability ε_f with $N = 20$.




Fig. 12. AoII violation as a function of the feedback deletion probability ω_f with $N = 20$.

most significant parameter settings, as well as to imperfect feedback and traffic load estimation errors. Furthermore, the performance gap widens as the number of nodes increases, making the protocol suitable for large sensor networks.

Our work also opens several possible extensions and research directions, from a case in which anomalies are modeled as a more complex N -state Markov process to a more complex case in which nodes have structured beliefs about their own and others' observations.

APPENDIX

DELTA CORRECTNESS ANALYSIS

We verify the correctness of the DELTA protocol by considering the system as an epistemic program, following the notation from [19]. We first define the Kripke state model [47] $\mathcal{S} = (S, \xrightarrow{\mathcal{N}}_S, \parallel \cdot \parallel_S)$. We can define a set of atomic propositions $\mathcal{P}(S)$ over the set of possible worlds S :

$$S = \{(\Delta, \Theta, \Phi) \in \mathbb{N}^{2N} \times \{ZW, CR, CE, BT\}^N : \Theta_n \leq \Delta_n \forall n \in \{1, \dots, N\}\}. \quad (36)$$

In other words, each state, or possible world, s represents a different combination of AoI and AoII values for the nodes, along with the DELTA phase Φ_n of each node. By definition, the AoII in any possible world is upper bounded by the AoI.

Each member $\xrightarrow{n}_S \subseteq S \times S$ of the family of accessibility relations $\xrightarrow{\mathcal{N}}_S$ expresses which states nodes n considers possible for a given real state, and the truth map $\parallel \cdot \parallel_S : \mathcal{P}(S)$ represents the truth of each possible atomic proposition about the state. We can also extend atomic propositions to sentences, whose truth value follows modal logic rules: we denote a generic sentence as φ , and the language of possible sentences as \mathcal{L} . If sentence φ is true in world s , we denote it as $s \models \varphi$.

Additionally, we define a knowledge operator $\square_M(\varphi)$, which specifies whether the subset of nodes $M \subseteq \mathcal{N}$ know $\varphi \in \mathcal{L}$, and a common knowledge operator $\square_M^*(\varphi)$, which follows the game theoretic definition: under common knowledge, even the fact that everyone knows φ is known to everyone, i.e., $\square_M^*(\varphi)$ requires not only $\square_M(\varphi)$ but also $\square_M(\square_M(\varphi))$, recursively. The pointed epistemic model (\mathcal{S}, s) includes both the real state of the world and its perception by the nodes.

We express the outcome of each communication slot with the set of possible *public announcements* \mathcal{A} [48]. The three possible outcomes, i.e., a silent slot, a successful transmission and a failure, are represented by symbols γ , σ , and ω , respectively. The possible announcements then include the final outcome of the transmission, as well as which nodes transmitted and which were silent, unless:

$$\mathcal{A} = (\cup_{n \in \mathcal{N}} (\sigma, \mathbf{u}_N^n)) \cup \{\gamma, \omega\}, \quad (37)$$

where \mathbf{u}_N^n is the column vector of length N whose elements are 0, except for element n , which is equal to 1. We denote the fact that statement φ holds after announcement α as $[\alpha!] \varphi$.

We then define four modes of operation. The first is the *default mode*, in which it is common knowledge that all nodes are in phase ZW and have AoII $\Theta_{n,t} = 0$, corresponding to proposition φ_d :

$$\varphi_d = \square_N^* ((\Theta_{n,t} = 0 \wedge \Phi_{n,t} = \text{ZW}) \forall n \in \mathcal{N}). \quad (38)$$

The *collision mode* and *exit mode* are then characterized by the common knowledge of the phase, which is CR for all nodes for the former and CE for all nodes for the latter. Additionally, the AoII vector ψ_t is also common knowledge. They correspond to propositions φ_c and φ_e , respectively:

$$\varphi_c = \square_N^* (\Phi_{n,t} = \text{CR} \forall n \in \mathcal{N}) \wedge \square_N^* (\psi_t); \quad (39)$$

$$\varphi_e = \square_N^* (\Phi_{n,t} = \text{CE} \forall n \in \mathcal{N}) \wedge \square_N^* (\psi_t). \quad (40)$$

Finally, the *threshold mode* is a state in which all nodes are in phase BT, and this is common knowledge, as is the maximum AoII vector ψ_t . However, the nodes do not have common knowledge about the actual AoII. This is expressed by φ_t :

$$\begin{aligned} \varphi_t = & \square_N^* (\psi_t \wedge (\Phi_{n,t} = \text{BT} \forall n \in \mathcal{N})) \\ & \wedge \neg \square_N^* (\Theta_{n,t} = 0 \forall n \in \mathcal{N}). \end{aligned} \quad (41)$$

We will now prove that, under an ideal feedback model, a DELTA network is always in one of these four modes.

Lemma 1.1. *If the network is in default mode in step t, there are two possible outcomes: either it remains in default mode in step t + 1, or it moves to collision mode. In DEL notation,*

$$(s_t \models \varphi_d) \Rightarrow (s_t \models \varphi_d \vee s_{t+1} \models \varphi_c). \quad (42)$$

Proof: Consider the three possible outcomes of slot t. If $\alpha_t = \gamma$, i.e., the slot is silent, no node transmitted, i.e., $a_{n,t} = 0 \forall n$. In the ZW phase, nodes are silent only if their state is normal, i.e., $a_{n,t} = 0 \Rightarrow x_{n,t} = 0$. We also have $x_{n,t} = 0 \Rightarrow \Theta_{n,t+1} = 0$, by definition. As the node remains in phase ZW after a silent slot, we have $(s_t \models \varphi_d) \Rightarrow [\omega!] (s_{t+1} \models \varphi_d)$. On the other hand, if $\alpha_t = (\sigma, n)$, we have $x_{n,t} = 1$, but the anomaly is immediately resolved, and $\Theta_{n,t+1} = 0$. On the other hand, any other node m is silent, so the proof follows the previous case. All nodes can follow this procedure, as the ACK contains the identity of n . We have $(s_t \models \varphi_d) \Rightarrow [(\sigma, n)!] (s_{t+1} \models \varphi_d)$. Finally, if $\alpha_t = \omega$, nodes move to the CR phase following a NACK. Since at least one node transmitted in slot t, we have $\exists n : a_{n,t} = 1 \Rightarrow \exists n : x_{n,t} = 1$. The maximum AoII $\psi_{n,t}$ is then equal to 1 for all nodes, as it is impossible to know n using only public information. We then have $(s_t \models \varphi_d) \Rightarrow [\omega!] (s_{t+1} \models \varphi_c)$. ■

The straightforward consequence of the lemma is that, whenever it is common knowledge that all nodes are in phase ZW, such knowledge is preserved until the switch to phase CR, which is performed simultaneously. While the knowledge of the actual AoII is lost, the maximum AoII vector is still common knowledge, and nodes can use it to coordinate.

Lemma 1.2. *If the network is in collision mode at step t, it will either remain in collision mode at step t + 1, or switch to exit mode. In DEL notation, we have*

$$(s_t \models \varphi_c) \Rightarrow (s_t \models \varphi_c \vee s_{t+1} \models \varphi_e). \quad (43)$$

Proof: We can follow the same steps as for the previous lemma: in this case, outcomes γ and ω are public announcements, and $(s_t \models \varphi_c) \Rightarrow [(\omega \vee \gamma)!] (s_{t+1} \models \varphi_c)$. The maximum AoII vector ψ_{t+1} is simply increased by 1, as the feedback contains no information on new or existing anomalies. On the other hand, since σ is also a public announcement, all nodes know that nodes move to phase CE when they receive an ACK. In this case, the maximum AoII vector ψ_{t+1} is increased by 1, except for the node n that transmitted, whose AoII is reset to 0. This can be computed by all nodes through the feedback, and $(s_t \models \varphi_c) \Rightarrow [(\sigma, n)!] (s_{t+1} \models \varphi_e) \forall n \in \mathcal{N}$. ■

Lemma 1.3. *If the network is in exit mode at step t, it will either go back to collision mode at step t + 1, or switch to threshold mode. In DEL notation, we have*

$$(s_t \models \varphi_e) \Rightarrow (s_t \models \varphi_t \vee s_{t+1} \models \varphi_c). \quad (44)$$

Proof: In the case of outcome ω , the DELTA protocol dictates that the nodes must move back to the CR phase, and $(s_t \models \varphi_e) \Rightarrow [\omega!] (s_{t+1} \models \varphi_c)$. Vector ψ_t is increased by 1 by all nodes. We can similarly prove that all nodes move to phase BT following the protocol if there is another outcome. If the outcome is a successful transmission, the maximum AoII for node n $\psi_{m,t+1} = \min(\psi_{m,t} + 1, k-1)$, as $\psi_{m,t} \geq k \Rightarrow x_{m,t-k} = 1$. If the state was 1 at time $t-k$, and $\Phi_{m,t-k} = \text{ZW}$, the node must have been part of the collision set: $\Phi_{m,t-k} = \text{ZW} \wedge x_{m,t-k} = 1 \Rightarrow a_{m,t-k} = 1$. This means that it was a member of the collision set, and must consequently have transmitted in the CE phase if its original anomaly is still unresolved, but this is a contradiction. Consequently, no node can have an AoII higher than $k-1$, and ψ_t is common knowledge. Finally, we also need to prove that it is not common knowledge that the AoII of all nodes is 0. In the best possible case, the collision is resolved in 2 steps, during which only members of the collision set can transmit. We consider AoII vector $\Theta_{t+1} = \mathbf{u}_N^m$, in which only node m , which was not in the collision set has an AoII equal to 1. We have $\Phi_{n,t} = \text{CE} \Rightarrow a_{m,t} = 0$ due to the rules of the CE phase. World $s'_{t+1} = (\Delta_{t+1}, \mathbf{u}_N^m, \mathbf{BT})$ is then compatible with the announcement, i.e., $(s_t \models \varphi_e) \Rightarrow [\omega!] (s'_{t+1} \models \varphi_c)$. However, this is in contradiction with the common knowledge proposition $\square_N^* (\Theta_{n,t} = 0 \forall n \in \mathcal{N})$, and φ_t must therefore be verified, proving the lemma. ■

Lemma 1.4. *If it is common knowledge that all nodes are always in the same phase, the maximum AoII vector ψ_t is common knowledge at any time:*

$$\square_N^* (\Phi_{m,t} = \Phi_{n,t}) \forall m, n \in \mathcal{N}, t \in \mathbb{N} \Rightarrow \square_N^* (\psi_t) \forall t \in \mathbb{N}. \quad (45)$$

Proof: We have $\varphi_d \Rightarrow \square_{\mathcal{N}}^*(\psi_{n,t} = 0 \ \forall n \in \mathcal{N})$ by definition. If we consider the sequence of CR and CE phases starting at time t from default mode and ending after k slots, there are two common knowledge propositions: firstly, as stated in Lemma 1.2, switches between these two phases are common knowledge. We know that $s_t \models \varphi_d \Rightarrow \psi_{n,t+k} \leq k$. Let us then consider the case in which a node m has $x_{m,t} = 0$ and $x_{m,t+1} = 1$. We know that $\Phi_{m,t+j} \in \{\text{CR, CE}\} \wedge x_{m,t} = 0 \Rightarrow a_{n,t+j} = 0 \ \forall j \in \{t+1, \dots, t+k\}$: as node m was not in the collision set, it cannot transmit until time $t+k$. There is then a world, compatible with the feedback history, in which $\psi_{m,t} = k+1$. The same reasoning can be applied to nodes in the collision set, as it is common knowledge that they reset their AoII at the time of their last transmission. During phase BT, we can easily compute $\psi_{n,t+1} = \sup\{\theta \in \mathbb{N} : z_{n,t}(\theta, \psi_t) \leq Z\}$, and any collisions during the BT phase can be managed in the same way. ■

Lemma 1.5. *If the network is in threshold mode at step t , it will remain in threshold mode at step $t+1$ or switch to either collision or default mode. In DEL notation, we have*

$$\square_{\mathcal{N}}^*(\psi_t) \wedge (s_t \models \varphi_t) \Rightarrow (s_t \models \varphi_t \vee s_{t+1} \models \varphi_c \vee s_{t+1} \models \varphi_d). \quad (46)$$

Proof: If the network is in threshold mode, we can analyze the effect of different outcomes. In case of a collision, all nodes will move to the CR phase. If ψ_t is common knowledge, it is also common knowledge that $\psi_{t+1} = \psi_t + 1$, and $s_t \models \varphi_t \Rightarrow [\omega!] (s_{t+1} \models \varphi_c)$. In case of a silent or successful slot, we can easily compute $\psi_{t+1} = \sup\{\theta \in \mathbb{N} : z_{n,t}(\theta, \psi_t) \leq Z\}$, except for the node that successfully transmitted, whose maximum AoII is 0. Consequently, $s_t \models \varphi_t \Rightarrow \square_{\mathcal{N}}^*(\psi_t + 1)$. If the maximum AoII is 0, the real AoII must also be 0: $\psi_{t+1} = 0_N \Rightarrow \Theta_t = 0_N$. In this case, all nodes will switch to ZW, while in other cases, they will remain in BT. We then have $s_t \models \varphi_t \Rightarrow [(\gamma \vee (\sigma, n))!] (s_{t+1} \models (\varphi_t \vee \varphi_d))$. ■

We can then restate Theorem 1 using DEL notation, proving its correctness using the above Lemmas.

Theorem 1. *The phase and the value of ψ_t are always common knowledge, and the phase is always the same for all nodes:*

$$\square_{\mathcal{N}}^*(\psi_{n,t}, \Phi_{n,t}) \ \forall n \in \mathcal{N} \wedge \square_{\mathcal{N}}^*(\Phi_{n,t} = \Phi_{m,t} \ \forall m, n \in \mathcal{N}). \quad (47)$$

Proof: Under an ideal feedback channel, the system is always in one of the four modes, i.e.,

$$s_t \models (\varphi_d \vee \varphi_c \vee \varphi_e \vee \varphi_t) \ \forall t \in \mathbb{N}, \quad (48)$$

due to the concatenation of the four preceding lemmas: as the system starts in default mode in step 0, it will remain in one of the four modes at all times, and this is common knowledge. As all modes require ψ_t to be common knowledge, this is an immediate consequence. Additionally, all modes require every node to have the same phase, proving the theorem. ■

REFERENCES

- [1] F. Chiariotti, A. Munari, L. Badia, and P. Popovski, “Distributed optimization of age of incorrect information with dynamic epistemic logic,” in *Proc. IEEE Comp. Commun. Conf. (INFOCOM)*, 2025.
- [2] D. Gündüz, Z. Qin, I. E. Aguerri, H. S. Dhillon *et al.*, “Beyond transmitting bits: Context, semantics, and task-oriented communications,” *IEEE J. Sel. Areas Commun.*, vol. 41, no. 1, pp. 5–41, 2023.
- [3] E. Bourtsoulatze, D. B. Kurka, and D. Gündüz, “Deep joint source-channel coding for wireless image transmission,” *IEEE Trans. on Cogn. Commun. Netw.*, vol. 5, no. 3, pp. 567–579, 2019.
- [4] Y. Shao, Q. Cao, and D. Gündüz, “A theory of semantic communication,” *IEEE Trans. Mobile Comput.*, vol. 23, no. 12, pp. 12211–12228, 2024.
- [5] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, “Minimizing age of information in vehicular networks,” in *Proc. IEEE Int. Conf. Sensing Commun. Netw. (SECON)*, 2011, pp. 350–358.
- [6] A. Kosta, N. Pappas, and V. Angelakis, “Age of information: A new concept, metric, and tool,” *Found. Trends Netw.*, vol. 12, no. 3, pp. 162–259, 2017.
- [7] Z. Lu, R. Li, K. Lu, X. Chen *et al.*, “Semantics-empowered communications: A tutorial-cum-survey,” *IEEE Commun. Surveys Tuts.*, vol. 26, no. 1, pp. 41–79, 2024.
- [8] 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.
- [9] C. De Alwis, A. Kalla, Q.-V. Pham, P. Kumar *et al.*, “Survey on 6G frontiers: Trends, applications, requirements, technologies and future research,” *IEEE Open J. Commun. Soc.*, vol. 2, pp. 836–886, 2021.
- [10] J. Holm, F. Chiariotti, A. E. Kalør, B. Soret *et al.*, “Goal-oriented scheduling in sensor networks with application timing awareness,” *IEEE Trans. Commun.*, vol. 71, no. 8, pp. 4513–4527, 2023.
- [11] O. Ayan, S. Hirche, A. Ephremides, and W. Kellerer, “Optimal finite horizon scheduling of wireless networked control systems,” *IEEE/ACM Trans. Netw.*, vol. 32, no. 2, pp. 927–942, 2024.
- [12] L. Badia, “Impact of transmission cost on age of information at Nash equilibrium in slotted ALOHA,” *IEEE Netw. Lett.*, vol. 4, no. 1, pp. 30–33, 2021.
- [13] A. Fahim, T. Elbatt, A. Mohamed, and A. Al-Ali, “Towards extended bit tracking for scalable and robust RFID tag identification systems,” *IEEE Access*, vol. 6, pp. 27190–27204, 2018.
- [14] G. C. Madueño, Č. Stefanović, and P. Popovski, “Efficient LTE access with collision resolution for massive m2m communications,” in *Proc. IEEE Globecom Workshops (GC Wkshps)*, 2014, pp. 1433–1438.
- [15] A. Nayak, A. E. Kalør, F. Chiariotti, and P. Popovski, “A decentralized policy for minimization of age of incorrect information in slotted ALOHA systems,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2023, pp. 1688–1693.
- [16] O. T. Yavascan and E. Uysal, “Analysis of slotted ALOHA with an age threshold,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1456–1470, 2021.
- [17] A. Munari, “Monitoring IoT sources over random access channels: Age of incorrect information and missed detection probability,” in *Proc. Int. Conf. Commun. (ICC)*, IEEE, 2024, pp. 207–213.
- [18] L. Li, Y. Dong, C. Pan, and P. Fan, “Timeliness of wireless sensor networks with random multiple access,” *IEEE/KICS J. Commun. Netw.*, vol. 25, no. 3, pp. 405–418, 2023.
- [19] A. Baltag and L. S. Moss, “Logics for epistemic programs,” *Synthese*, vol. 139, pp. 165–224, 2004.
- [20] A. Maatouk, M. Assaad, and A. Ephremides, “On the age of information in a CSMA environment,” *IEEE/ACM Trans. Netw.*, vol. 28, no. 2, pp. 818–831, 2020.
- [21] H. Pan, T.-T. Chan, J. Li, and V. C. Leung, “Age of information with collision-resolution random access,” *IEEE Trans. Veh. Technol.*, vol. 71, no. 10, pp. 11295–11300, 2022.
- [22] M. A. Abd-Elmagid, N. Pappas, and H. S. Dhillon, “On the role of age of information in the Internet of Things,” *IEEE Commun. Mag.*, vol. 57, no. 12, pp. 72–77, 2019.
- [23] J. Yang, D. B. Minturn, and F. T. Hady, “When poll is better than interrupt,” in *Proc. Conf. File Storage Techn. (FAST)*, vol. 12. USENIX, 2012.
- [24] L. Scheuvens, T. Hößler, P. Schulz, N. Franchi *et al.*, “State-aware resource allocation for wireless closed-loop control systems,” *IEEE Trans. Commun.*, vol. 69, no. 10, pp. 6604–6619, 2021.
- [25] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul *et al.*, “Age of information: An introduction and survey,” *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, 2021.
- [26] A. M. Bedewy, Y. Sun, S. Kompella, and N. B. Shroff, “Optimal sampling and scheduling for timely status updates in multi-source networks,” *IEEE Trans. Inf. Theory*, vol. 67, no. 6, pp. 4019–4034, 2021.
- [27] A. Munari and E. Uysal, “Information freshness in random access channels for IoT systems,” in *Proc. IEEE Int. Balkan Conf. Commun. Netw. (BalkanCom)*, 2021, pp. 112–116.

- [28] M. Moradian, A. Dadlani, A. Khonsari, and H. Tabassum, "Age-aware dynamic frame slotted ALOHA for machine-type communications," *IEEE Trans. Commun.*, vol. 72, no. 5, pp. 2639–2654, 2024.
- [29] W. De Sombre, F. Marques, F. Pyttel, A. Ortiz, and A. Klein, "A unified approach to learn transmission strategies using age-based metrics in point-to-point wireless communication," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, 2023, pp. 3573–3578.
- [30] Z. Jiang, B. Krishnamachari, X. Zheng, S. Zhou, and Z. Niu, "Timely status update in wireless uplinks: Analytical solutions with asymptotic optimality," *IEEE Internet Things J.*, vol. 6, no. 2, pp. 3885–3898, 2019.
- [31] J. Thota, N. F. Abdullah, A. Doufexi, and S. Armour, "V2V for vehicular safety applications," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 6, pp. 2571–2585, 2020.
- [32] R. Kontar, N. Shi, X. Yue, S. Chung, E. Byon, M. Chowdhury, J. Jin, W. Kontar, N. Masoud, M. Nouiehed *et al.*, "The Internet of Federated Things (IoFT)," *IEEE Access*, vol. 9, pp. 156 071–156 113, 2021.
- [33] Z. Ning, P. Dong, X. Wang, X. Hu *et al.*, "Mobile edge computing enabled 5G health monitoring for Internet of medical things: A decentralized game theoretic approach," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 2, pp. 463–478, 2021.
- [34] E. Uhlemann, "Connected-vehicles applications are emerging," *IEEE Veh. Technol. Mag.*, vol. 11, no. 1, pp. 25–96, 2016.
- [35] H. Hui, S. Wei, and W. Chen, "Fresh multiple access: A unified framework based on large models and mean-field approximations," *IEEE/KICS J. Commun. Netw.*, vol. 25, no. 5, pp. 598–620, 2023.
- [36] G. Cocco, A. Munari, and G. Liva, "Remote monitoring of two-state Markov sources via random access channels: an information freshness vs. state estimation entropy perspective," *IEEE J. Sel. Areas Inf. Th.*, vol. 4, pp. 651–666, 2023.
- [37] F. Malandrino, C. F. Chiasserini, and G. Di Giacomo, "Efficient distributed DNNs in the mobile-edge-cloud continuum," *IEEE/ACM Trans. Netw.*, vol. 31, no. 4, pp. 1702–1716, 2023.
- [38] L. Huang, Y. Wu, J. M. Parra-Ullauri, R. Nejabati, and D. Simeonidou, "AI model placement for 6G networks under epistemic uncertainty estimation," in *Proc. Int. Conf. Commun. (ICC)*. IEEE, 2024, pp. 5521–5526.
- [39] X. Chen and H. Deng, "Analysis of cryptographic protocol by dynamic epistemic logic," *IEEE Access*, vol. 7, pp. 29 981–29 988, 2019.
- [40] J. J. Kline, "Evaluations of epistemic components for resolving the muddy children puzzle," *Econ. Th.*, vol. 53, no. 1, pp. 61–83, 2013.
- [41] L. Badia, "On the effect of feedback errors in Markov models for SR ARQ packet delays," in *Proc. Global Commun. Conf. (Globecom)*, 2009.
- [42] "Study on low-power wake-up signal and receiver for NR," 3GPP, Tech. Rep., 2023.
- [43] N. Mazloum and O. Edfors, "Interference-free OFDM embedding of wake-up signals for low-power wake-up receivers," *IEEE Trans. on Green Commun. Netw.*, vol. 4, no. 3, pp. 669–677, 2020.
- [44] A. Höglund, M. Mozaffari, Y. Yang, G. Moschetti, K. Kittichokechai, and R. Nory, "3GPP release 18 wake-up receiver: Feature overview and evaluations," *IEEE Commun. Standards Mag.*, vol. 8, no. 3, pp. 10–16, 2024.
- [45] K. Sikorski, "Bisection is optimal," *Numerische Mathematik*, vol. 40, no. 1, pp. 111–117, 1982.
- [46] Y. C. Tay, K. Jamieson, and H. Balakrishnan, "Collision-minimizing CSMA and its applications to wireless sensor networks," *IEEE J. Sel. Areas Commun.*, vol. 22, no. 6, pp. 1048–1057, 2004.
- [47] S. Kripke, "Semantical considerations on modal logic," *Acta Phil. Fennica*, vol. 16, pp. 83–94, 1963.
- [48] Y. Wang and Q. Cao, "On axiomatizations of public announcement logic," *Synthese*, vol. 190-Suppl. 1, pp. 103–134, 2013.



Federico Chiariotti (Senior Member, IEEE) is currently an Assistant Professor at the Department of Information Engineering, University of Padova, Italy, where he also received his PhD in 2019. From 2020 to 2022, he joined Aalborg University, Denmark. His research focuses on Age and Value of Information and goal-oriented communications, and he has published over 100 peer-reviewed papers, receiving 4 Best Paper Awards. He is an Associate Editor for the IEEE Networking Letters and IEEE Transactions on Wireless Communications.



Andrea Munari (Senior Member, IEEE) received the Ph.D. in information engineering from the University of Padova, Padua, Italy, in 2010. From 2007 to 2010, he was with IBM Research in Zurich, Switzerland. In 2011, he joined the Corporation Research and Development Division, Qualcomm Inc., San Diego, CA, US. He is currently with the Institute of Communications and Navigation, German Aerospace Center (DLR), Wessling, Germany. From 2014 to 2018, he held a Senior Researcher and a Lecturer position at the RWTH Aachen University, Germany. His main research interests are in the area of wireless communications, with special focus on Internet of things and satellite communications. Dr. Munari received the 2018 ACM MobiCom Workshop on Millimeter Wave Networks and Sensing Systems Best Paper Award and the IEEE GLOBECOM 2020 Communications Theory Symposium Best Paper Award. He serves as Associate Editor of IEEE Communications Letters.



Leonardo Badia (Senior Member, IEEE) received the Ph.D. degree in Information Engineering from the University of Ferrara, Italy, in 2004. From 2002 to 2003, he was a Visiting Scholar at the Radio System Technology Labs (currently Wireless@KTH), Royal Institute of Technology, Stockholm, Sweden. In 2006, he joined the IMT Institute for Advanced Studies, Lucca, Italy. In 2011, he moved to the University of Padova, Italy, where he is currently a Full Professor. He has published more than 300 research papers. His research interests include cross-layer optimization in wireless communications, transmission protocol modeling, and the application of game theory to radio resource management. He has also served on the editorial boards of several scientific journals and remains an active reviewer for many international periodicals, as well as a technical program committee member for conferences in the broad areas of communications and networking.



Petar Popovski (Fellow, IEEE) is a Professor at Aalborg University, where he is the director of the Center of Excellence CLASSIQUE (Classical Communication in the Quantum Era). He also holds a position of a Visiting Excellence Chair at the University of Bremen. He received his Dipl.-Ing (1997) and M. Sc. (2000) degrees in communication engineering from the University of Sts. Cyril and Methodius in Skopje and the Ph.D. degree (2005) from Aalborg University. He is a Fellow of the IEEE. He received an ERC Consolidator Grant (2015), the Danish Elite Researcher award (2016), IEEE Fred W. Ellersick prize (2016), IEEE Stephen O. Rice prize (2018), Technical Achievement Award from the IEEE Technical Committee on Smart Grid Communications (2019), the Danish Telecommunication Prize (2020) and Villum Investigator Grant (2021). He was a Member at Large at the Board of Governors in IEEE Communication Society, Chair of the IEEE Communication Theory Technical Committee, as well as the General Chair for IEEE SmartGridComm 2018 and IEEE Communication Theory Workshop 2019. His research interests are in the area of wireless communication and communication theory. He authored the book "Wireless Connectivity: An Intuitive and Fundamental Guide", published by Wiley in 2020. He is currently an Editor-in-Chief of IEEE Journal on Selected Areas in Communications.