

Distributed Optimization of Age of Incorrect Information with Dynamic Epistemic Logic

Federico Chiariotti*, Andrea Munari[†], Leonardo Badia*, and Petar Popovski[‡]

*Department of Information Engineering, University of Padova, 35131 Padua, Italy

[†]Institute of Communications and Navigation, German Aerospace Center (DLR), Weßling, Germany

[‡]Department of Electronic Systems, Aalborg University, 9220 Aalborg, Denmark

Emails: chiariot@dei.unipd.it, andrea.munari@dlr.de, leonardo.badia@unipd.it, petarp@es.aau.dk

Abstract—Distributed medium access schemes have a key advantage in anomaly tracking applications, as individual sensors know their own observations and can exploit them to reduce their Age of Incorrect Information (AoII). However, the risk of collisions has so far limited their performance. 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. This is achieved by a process of inferring AoII from plain Age of Information (AoI). In a network scenario with randomly generated anomalies, the individual AoII for each sensor is known only to itself, but all nodes can infer its AoI by simply tracking the transmission process. Thus, we adopt an approach based on dynamic epistemic logic, which allows individual nodes to infer how their AoII values rank among the entire network by exploiting public information such as the AoI and the identity of transmitting nodes. We analyze the resulting DELTA protocol both from a theoretical standpoint and with Monte Carlo simulation, showing that our approach is significantly more efficient and robust than basic random access, while outperforming state-of-the-art scheduled schemes by at least 30%.

Index Terms—Age of incorrect information; age of information; dynamic epistemic logic; medium access control.

I. INTRODUCTION

Future networks will support a plethora of real-time applications, entailing the exchange of timely and accurate status information among sensing and control units [1]. Age of Information (AoI), which represents the time elapsed since the generation of the last received status report, has emerged as a popular, amenable to quantitative analysis [2], [3] metric for information freshness. Extensions of AoI supplement the timing with the data content of the process. Specifically, Age of Incorrect Information (AoII) considers a linear penalty counting the time elapsed since the last variation of system conditions [4], thereby supplementing staleness with the semantic character of anomaly tracking [5].

Besides the basic single-link context for which AoI and AoII have been defined, there is a growing interest in setups where AoI and AoII are monitored and optimized for multiple

agents that share the communication medium. Most schemes consider centralized setups, which avoid collisions by having the receiving gateway coordinate transmissions [6]. This has some significant drawbacks in Internet of Things (IoT) settings: (i) dynamic scheduling requires nodes to listen to polling requests, causing resource reservation problems, and (ii) the large number of sensors and relative rarity of anomalous events can increase latency. This points towards protocols that use random access to share the medium among uncoordinated agents, as in ALOHA [7]. More advanced protocols based on splitting trees and collision resolution [8] require that contending nodes use the feedback from the common receiver to compute the transmission state of other contending nodes. More in general, random access protocols need to be redesigned when the objective is information freshness and AoII [9]–[11] rather than, traditionally, throughput.

This sets the motivation for the present work. We consider random access for a setup where the objective is to optimize information freshness with respect to the instant at which a contending node enters into a specific state; e.g., occurrence of anomaly. This model can be illustrated by the following two applications. (i) A set of wireless sensors reports to a common access point. Upon the occurrence of an anomaly, e.g., a wildfire, the sensor that detects it enters an alert state. The objective is, then, to minimize the time required to report the anomaly [12]. (ii) A scenario in which agents request access to computing resources over a shared channel. Upon getting a computing task, a node tries persistently to send a request/interrupt to the common computing engine [13] and get computing time. Data transmission between the node and the computing engine can happen over a different, dedicated channel [14]; what we are concerned with here is how fast the computing request can get through a shared channel [15].

The formal framework that represents these two applications is the optimization of the AoII. We look for an efficient access control that can still exploit the sensors' knowledge of their measurements, while avoiding the pitfalls of random access schemes [16]. We find the answer in deductive reasoning over the information freshness. The nodes are aware of their own AoII but not that of the others; the key observation is that they can obtain the plain AoI of every node by observing the transmission patterns and listening to broadcast feedback packets. Applying Dynamic Epistemic Logic (DEL) [17], they

This work was supported in part by the Italian National Recovery and Resilience Plan (NRRP), as part of the RESTART partnership (PE00000001) and the “Young Researchers” grant REDIAL (SoE00000009), under the European Union NextGenerationEU Project, by the Federal Ministry of Education and Research of Germany in the programme of “Souverän. Digital. Vernetzt.” joint project 6G-RIC, project identification number: 16KISK022, and by the Villum Investigator Grant “WATER” from the Velux Foundations, Denmark.

can therefore use *common knowledge* information to coordinate behavior, thereby reducing the occurrence of collisions and shortening the collision resolution phase [18].

1) *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 that is extremely sporadic [11], [19]. Even though it was studied in the seminal paper that first defined AoI [2], 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 [20], [21], 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. Several recent studies [10] considering AoI in random access channels point out how collisions have a detrimental effect on AoI, even when considering carrier sensing [16] and collision resolution mechanisms [18]. The efforts to prevent nodes from entering collisions are mostly circumscribed to the threshold ALOHA approach [9], which can be adapted dynamically to time-varying traffic conditions [22]. 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 [23].

Deterministic access quickly becomes AoI-optimal for large networks [24]; 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 [21]. Scenarios include vehicular flow management in which critical reporting by a vehicle is not constant and depends on its position [25], environmental supervision in smart agriculture, wildlife tracking, or safety and security monitoring in domotic, industrial, or smart grid scenarios [26]. Even medical supervision of elderly or chronic patients likely only reports relevant condition changes [27]. 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 [28]), 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 [29], while precise results for finite networks have been provided just for simple schemes, such as ALOHA [30], [31]. To the best of our knowledge, the only work to directly optimize the AoII instead of analyzing existing schemes is [32]; yet, performance is inferior to simple round-robin.

2) *Epistemic logic*: Epistemic logic is a branch of formal reasoning dealing with the inference, transfer, and update of knowledge among multiple agents [33], [34]. When knowledge evolves over time and successive interactions, this is referred to as DEL, and finds applications in social networks and

cryptography [35]. 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 [36]. 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, [26] 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, [34] 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.

3) *Contribution*: We design Dynamic Epistemic Logic for Tracking Anomalies (DELTA), a protocol that adopts epistemic logic to allow sensors to minimize AoII distributedly. Each node estimates the probability of its update being the one with the highest AoII. Listening to acknowledgments guarantees that each user is able to track everybody else’s AoI, and in turn this can be used to infer the likelihood of being the node with highest AoII and apply a threshold criterion to it. This method does not entirely prevent collisions, but is able to significantly mitigate them. At the same time, it allows for faster collision resolution, since in the event of simultaneous transmissions, i.e., whenever more than one user estimates their AoII to be the highest, they can reevaluate their belief in light of the new information gained. DELTA is specifically meant for tracking anomalies, or in general any kind of events that nodes should report as soon as possible to meet freshness constraints, but an aggressive medium access would cause too many collisions, thereby increasing staleness.

Thus, our approach presents the following novelties with respect to the corpus of existing literature. First, we exploit epistemic logic, which is already quite rare in the field. To the best of our knowledge, we are the first to do so from the perspective of the individual nodes, rather than a central network entity with a holistic view. Moreover, we design a random access protocol for anomaly reporting that not only explores an uncharted territory but is also shown to obtain superior performance over state-of-the-art approaches.

The contributions of this paper are listed as follows. (i) We formally prove that our proposed random access protocol, based on inference reasoning, can lead multiple sensors to efficiently operate based on common knowledge information. (ii) 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. (iii) We show how DELTA outperforms both scheduled and random access legacy protocols. This is proven both formally and through extensive Monte Carlo simulations, showing that sensors can operate based




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

on an inference logic from common knowledge information. (iv) In addition to the version of DELTA designed under the assumption of a perfect feedback, we discuss the design of a robust version that works with imperfect feedback. 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. The robust version exhibits graceful performance degradation under imperfect feedback.

II. 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 $\mathbf{x}_t \in \{0, 1\}^N$. At any time slot, sensor n may switch from the normal state 0 to the anomalous state 1 with probability λ_n . On the other hand, state 1 is absorbing, i.e., the anomaly persists until the sensor successfully transmits a warning to the gateway. The transition matrix $\mathbf{A}_n$ is then

$$\mathbf{A}_n = \begin{pmatrix} 1 - \lambda_n & \lambda_n \\ s_{n,t} & 1 - s_{n,t} \end{pmatrix}, \quad (1)$$

where $s_{n,t} \in \{0, 1\}$ is an indicator variable which is equal to 1 if n successfully transmits at time t and 0 otherwise.¹ We can then define the AoI $\Delta_{n,t}$ as

$$\Delta_{n,t} = t - \max_{\delta \in \{1, \dots, t\}} \delta s_{n,t-\delta}. \quad (2)$$

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 $\Theta_{n,t}$ [4], which is defined as

$$\Theta_{n,t} = t - \arg \max_{\theta \in \{t - \Delta_{n,t} + 1, \dots, t\}} \theta x_{n,t-\theta}. \quad (3)$$

As Fig. 1 shows, the AoI grows even while in the normal state, while the AoII only grows in the anomalous state. Our objective is then to minimize the AoII violation probability for a given AoII threshold, $V(\theta_{\text{thr}})$.

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 acknowledgment (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.

Most of the following assumes that the gateway transmits ACK and NACK packets with enough power to ensure their correct reception, as a shared feedback signal is crucial to maintain common knowledge. However, the protocol is robust to an imperfect feedback channel, and we will discuss the countermeasures to deal with a feedback erasure probability $\varepsilon_f > 0$, considering independent failures among sensors.

III. THE DELTA PROTOCOL

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 [9], [11], [16]. 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 often causes inefficiencies that have made distributed protocols valuable only for niche applications: to reduce the risk of collisions, sensors need to randomly abstain from transmitting, increasing their AoII even when there would be no need to do so.

The Dynamic Epistemic Logic for Tracking Anomalies (DELTA) protocol is based on the notion of *common knowledge* as defined in [17]. 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 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. 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 will consider the feedback channel to be ideal, but we will discuss how to adapt DELTA to an imperfect feedback channel in Sec. III-D.

A. Protocol Definition and Correctness

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 *Zero-Wait* (ZW) phase is the normal state of operation: during this phase,

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.

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 NACK signal to inform sensors of the collision, all sensors switch to the *Collision Resolution (CR)* phase [18], recording their membership in the collision set through an indicator variable $m_{n,t}$.

Lemma 1.1. *As long as the system remains in the ZW phase, all sensors are in state 0, and this is common knowledge.*

Proof: We can prove the lemma by induction. Let us consider slot t , knowing that all sensors are in state 0 at time $t - 1$. Since nodes in state 1 always transmit, a silent slot, in which case nobody had anything to transmit, can be interpreted by all nodes as the state remaining the same [37]. The same holds for a successful transmission, i.e., a single node transmitting and resetting its AoII and state to 0. On the other hand, a NACK may be caused by a wireless channel loss or a collision, moving all nodes to the CR phase. ■

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 increases the resolution time if there are more than 2 colliding nodes, as it causes another collision, but this case is relatively rare due to the low traffic, and it confers a major advantage: the second collision allows all nodes to know that the initial collision is still unresolved, and that there should be another CR phase. Conversely, successful or silent slots only happen when the collision set becomes empty, and nodes can safely switch from the CE to the *Belief Threshold (BT)* phase.

Lemma 1.2. *The switches between phases CR, CE, and BT are common knowledge if the feedback channel is ideal.*

Proof: After the switch from ZW to CR, state $\mathbf{x}_t$ is not common knowledge any more: each node knows its own state and AoII, but not others'. However, we can use public announcements to infer phase changes: if a transmission in the CR phase is successful, the transmitting node was part of the collision set, but its state is reset to 0, and the system switches to CE. The reception of an ACK in the CR phase then triggers to switch to the CE phase, and we note that ACKs are received by every sensor. We can make the same argument for the CE phase: as all remaining members of the collision set transmit, we know that the set is non-empty only after a NACK. All sensors then know that everyone switches back to CR after receiving a NACK in the CE phase, or to BT otherwise. ■

Finally, the BT phase allows sensors to gradually go back to




Fig. 2: DELTA state diagram.

normal: 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 $\psi_{n,t}$. By definition, $\Theta_{n,t} \leq \psi_{n,t}$, $\forall m, n$. Node n 's AoII $\Theta_{n,t}$ is the highest if no node has a 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

$$f_{n,t}(\Theta_{n,t}, \psi_t) = \prod_{m \neq n} (1 - \lambda_m)^{[\psi_{m,t} - \Theta_{n,t} + 1]^+}, \quad (4)$$

where $[x]^+ = \max(0, x)$ is the positive part operator. In the BT phase, we set a threshold F , and node n transmits with probability 1 if $f_{n,t} > F$. If $\psi_t = \mathbf{0}_N$, i.e., the vector of length N whose elements are all 0, the system goes back to the ZW phase. The DELTA phase diagram is shown in Fig. 2.

Theorem 1. *The protocol phase and ψ_t are always common knowledge if the feedback channel is ideal.*

Proof: As a direct consequence of Lemma 1.1, $\psi_{n,t} = 0$, $\forall n \in \mathcal{N}$ during the ZW phase. If we consider the sequence of CR and CE phases starting at time t from phase ZW and ending after k slots, there are two common knowledge propositions: firstly, as stated in Lemma 1.2, switches between phases are common knowledge. Secondly, it is common knowledge that nodes outside the collision set were in state 0 at time t , as they were in the ZW phase at that time and did not transmit.

The nodes with an AoI lower than j were in the collision set, and their transmissions reset their state to 0: their AoII is capped to their AoI by definition. When the BT phase begins,

$$\psi_{n,t+k} = \min(k, \Delta_{n,t}) \quad \forall n \in \mathcal{N}. \quad (5)$$

During the BT phase, communication decisions are based on the probability defined in (4). The outcome of each slot is then broadcasted: if sensor n did not transmit at time t ,

$$\psi_{n,t+1} = 1 + \sup(\tilde{\theta} \in \{0, \dots, \psi_{n,t}\} : f_{n,t}(\tilde{\theta}, \psi_t) < F). \quad (6)$$

If the outcome was silence or a successful transmission, all nodes (except the successful one, whose AoII was reset to 0) were silent. On the other hand, if the outcome of the round was a collision, all nodes except the members of the collision set were silent, by definition. The value of $\psi_{n,t+1}$ can then safely be reset for all nodes, as all colliding nodes will transmit again before the next BT phase. During subsequent collision resolution cycles, $\psi_{n,t}$ just increases by the duration of the cycle. The return to phase ZW depends only on ψ_t . ■

Algorithm 1 Pseudocode of the DELTA protocol

```

Require: phase,  $F$ ,  $\mathbf{p}$ ,  $x_{n,t}$ , 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:    $\psi_t \leftarrow \text{UPDATE\_MAXIMUM\_POSSIBLE\_AOII}(\psi_{t-1})$ 
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  $m_{n,t} \times \text{BERNOULLI\_SAMPLE}(p(c_t))$ 
19:     case CE: return  $m_{n,t}$ 
20:     case BT: return  $\text{HIGHEST\_AOII\_PROB}(\theta_t, \psi_t) > F$ 

```

We note that collisions are more common in the BT phase than in the ZW phase, as the threshold must necessarily be lower to gradually reduce ψ_t . Collisions occurring during the BT phase are handled identically to the ones in the ZW phase. The full decision-making algorithm and state update for each sensor is presented as Algorithm 1.

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 resolved. The probability of success in any given slot is

$$\sigma(c, p, \varepsilon) = (1 - \varepsilon_n) \text{Bin}(1; c, p), \quad (7)$$

where $\text{Bin}(k; N, p) = \binom{N}{k} p^k (1 - p)^{N-k}$ is the binomial Probability Mass Function (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

$$\mathbf{P}_c = \begin{pmatrix} \mathbf{B} & \sigma(2, c-1) \mathbf{u}_{c-1}^{c-1} \\ (\mathbf{0}_{c-1})^\top & 1 \end{pmatrix}, \quad (8)$$

where $\mathbf{u}_N^n$ is identical to $\mathbf{0}_N$ except for element n , which is equal to 1, and the elements of matrix $\mathbf{B}$ are²

$$B_{ij} = \begin{cases} 1 - \sigma(c - i + 1, p_i, \varepsilon), & j = i; \\ \sigma(c - i + 1, p_i, \varepsilon), & j = i + 1. \end{cases} \quad (9)$$

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 Cumulative Distribution Function

(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., a wireless channel error, 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, the expected duration of the subsequent CR and CE cycle is $\mathbb{E}[\tau_1] = 1 + ((1 - \varepsilon)p_1)^{-1}$. 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 after the transition to the BT phase, is*

$$\mathbb{E}[\tau_c] = c - 1 + \varepsilon + \frac{\varepsilon}{(1 - \varepsilon)p_c} + \sum_{i=0}^{c-2} \frac{1}{\sigma(c - i, p_{i+1}, \varepsilon)}. \quad (10)$$

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, computing the expected time until absorption of a Markov chain is complex, but the structure of the transition matrix allows us to simplify the calculation. Any state i is eventually reached from state $i - 1$ with a successful transmission after a geometrically distributed number of failures, i.e., self-transitions:

$$\mathbb{E}[\tau_{i-1,i} | C = c, p_{i-1}] = (\sigma(c - i, p_{i-1}, \varepsilon))^{-1}. \quad (11)$$

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, we have one more CE phase from the last colliding node when we have reached state c . If the transmission is successful, the cycle is over, but if there is a wireless channel loss, we have 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

$$p_C(c | \text{ZW}) = \text{Bin}(c; N, \lambda) [1 - (1 - \varepsilon)I(c = 1)], \quad (12)$$

where $I(\cdot)$ is the indicator function, equal to 1 if the condition in the argument is true and 0 otherwise. We can also easily get the total failure probability $p_f(\text{ZW}) = \sum_{c=1}^N p_C(c | \text{ZW})$. 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:

$$P_\tau(t | \text{ZW}) = \frac{\varepsilon(1 - \varepsilon)}{p_f(\text{ZW})} \left[\text{Bin}(1; N, \lambda) (1 - \eta_1^{t-1}) + \sum_{c=2}^{\min(N,t)} \text{Bin}(c; N, \lambda) \left(\frac{(\mathbf{P}_c)_{1,c}^{t-c+1}}{\varepsilon} + \sum_{k=1}^{t-2c+1} (\mathbf{P}_c)_{1,c}^{t-c-k} \eta_c^{k-1} p_c \right) \right], \quad (13)$$

where $\eta_c = 1 - (1 - \varepsilon)p_c$.

²We omit transitions with probability 0 for the sake of brevity.

Theorem 3. *There is a single optimal transmission strategy $\mathbf{p}^*$ that minimizes the expected duration*

$$\mathbf{p}^* = \arg \min_{\mathbf{p} \in (0,1)^N} \sum_{c=1}^N p_C(c|\mathbf{Z}\mathbf{W}) \mathbb{E}[\tau_c], \quad (14)$$

if all nodes have the same λ and ε , and p_i^* is the solution of

$$\text{Bin}(1; N_i, \lambda) \varepsilon + \sum_{c=2}^{N_i} \text{Bin}(c; N_i, \lambda) \frac{1 - cp_i^*}{c(p_i^*)^2(1 - p_i^*)^c} = 0, \quad (15)$$

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, we can optimize each element of $\mathbf{p}$ separately to minimize the expected duration; we can trivially see that this is equivalent to choosing the value of p_i that maximizes $\sigma(c - i + 1, p_i, \varepsilon)$. We then take the first probability:

$$\begin{aligned} p_1^* &= \arg \max_{p \in (0,1)} \left[\sum_{c=1}^N \frac{p_C(c|\mathbf{Z}\mathbf{W})(1 - (1 - \varepsilon)I(c=1))}{p_f(\mathbf{Z}\mathbf{W})\sigma(c, p, \varepsilon)} \right] \\ &= \arg \max_{p \in (0,1)} \left[\sum_{c=1}^N w_c \frac{1 - \varepsilon}{cp(1 - p)^{c-1}} \right]. \end{aligned} \quad (16)$$

In order to prove that it is convex, we only need to prove that each individual component is convex. The first one, with $c = 1$, is a simple linear function of p . To show that the components with $c > 1$ are also convex, we take the second derivative of $(\sigma(c, p, \varepsilon))^{-1}$ with respect to p :

$$\frac{\partial^2 (\sigma(c, p, \varepsilon))^{-1}}{\partial p^2} = (1 - \varepsilon) \frac{c(c+1)p^2 - 2(c+1)p + 2}{cp^3(1 - p)^{c+1}}. \quad (17)$$

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

$$c(c+1)p^2 - 2(c+1)p + 2 > 0. \quad (18)$$

This quadratic equation has no real solution for $c > 1$. 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. The maximum is then inside the interval for $N > 1$.

Finally, we can prove that (15) is a multiple of the first derivative of the optimization function in (14), and finding its root in $(0, 1)$ is equivalent to finding the minimum. As the solution of (15) involves a hypergeometric function, there is no closed-form solution, but it can be approximated efficiently with the bisection method and stored in a look-up table. ■

C. Belief Threshold Optimization

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




Fig. 3: Approximated semi-Markov model of the system with $K = 3N$ and $\Psi = 4$.

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 AoI ψ , we can obtain the conditioned PMF of the number of colliders by applying Bayes' theorem:

$$p_C(C|\psi) = (p_\tau(\psi))^{-1} p_C(c|\mathbf{Z}\mathbf{W}) p_{\tau_c}(\psi^*), \quad (19)$$

where $p_\tau(\psi)$ is the PMF corresponding to the CDF in (13).

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 $\psi_{n,t}$, while 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. The number of active nodes in the optimistic model is $L(\psi) = N - \mathbb{E}[C|\psi]$. Each sensor transmits with probability $\alpha = 1 - (1 - \lambda)^{\frac{K}{L(\psi)}}$, so the collision probability is

$$\xi(\psi) = 1 - (1 - \lambda)^K - (1 - \varepsilon) \text{Bin}(1; L(\psi), \alpha). \quad (20)$$

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 AoI Ψ , so that $|\mathcal{S}| = 2\Psi + 1$. We can reduce the approximation error as much as possible by considering 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 $\text{CR}(\Psi)$ does not

decrease as Ψ increases, the system is unstable. We can then give the elements of the transition matrix $\mathbf{M}$ of our model, considering the transitions toward state ZW:

$$M_{s,ZW} = (1 - \xi(\psi))I(s = \text{BT}(\psi))I(\psi L(\psi) < K). \quad (21)$$

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

$$M_{s,\text{CR}(\psi)} = \begin{cases} 1, & s = \text{ZW}, \psi = 0; \\ \xi(\psi'), & s = \text{BT}(\psi'), \psi' = \left[\psi + 1 - \frac{K}{L(\psi')} \right]^+. \end{cases} \quad (22)$$

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

$$M_{s,\text{BT}(\psi)} = \begin{cases} \zeta_{\psi'}(\psi - \psi'), & s = \text{CR}(\psi'); \\ 1 - \xi(\psi'), & s = \text{BT}(\psi + 1 - \frac{K}{L(\psi')}); \\ \sum_{\ell=\Psi-\psi'}^{\infty} \zeta_{\psi'}(\ell), & s = \text{CR}(\psi'), \psi = \Psi; \end{cases} \quad (23)$$

where $\zeta_{\psi'}(\ell)$ is the PMF corresponding to the CDF given in (13), 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 $T_{\text{ZW},\text{CR}(0)} = \text{Geo}(\xi(0))$, $T(\text{BT}(\psi), s') = 1$, and $T(\text{CR}(\psi), \text{BT}(\psi')) = \psi' - \psi$. We also consider a pessimistic approximation: if the collision resolution process leads to state $\text{BT}(\Psi)$, 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 α as the solution to the equation $\alpha(\mathbf{P} - \mathbf{I}) = 0$, normalized so that $\|\alpha\|_1 = 1$. This corresponds to the left eigenvector of $\mathbf{M}$ with eigenvalue 1. The steady-state distribution π is obtained by weighting α by the average sojourn times $\mathbb{E}[T(s, s')]$:

$$\pi(s) = \frac{\sum_{s' \in \mathcal{S}} \alpha(s') M(s, s') \mathbb{E}[T(s, s')]}{\sum_{s^*, s^{**} \in \mathcal{S}} \alpha(s^*) M(s^*, s^{**}) \mathbb{E}[T(s^*, s^{**})]}, \quad \forall s \in \mathcal{S}. \quad (24)$$

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

D. Dealing with Imperfect Feedback

If the feedback channel is imperfect, Theorem 1 does not hold. To compute ψ_t and synchronize phase transitions, all nodes need to receive 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




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

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 $f_{n,t}(\Theta_{n,t}, \psi_t)$ 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 acknowledgment packet, 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. Furthermore, proving its correctness might require an automated formal verification, as the DEL model considerably increases with respect to the ideal feedback case. We will leave this extension as future work.

IV. SIMULATION SETTINGS AND RESULTS

This section presents the results of the Monte Carlo simulations meant to validate the performance of the DELTA protocol. Each considered setting was tested over a simulation lasting 10^7 slots, aside from the benchmark protocol optimization, which used 10^6 slots for each setting due to the significant number of runs required by the grid search. In the following, the maximum offered system load $\rho = \|\lambda\|_1^3$ will be considered as the main simulation parameter.⁴

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 $\pi(\text{ZW})$, 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

³As the generation of new anomalies depends on the resolution time, this is an upper bound to the actual load experienced by the system.

⁴The complete 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 .

of collisions in this case: most of the time, higher values of ψ will be reached due to multiple collisions among 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 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.

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 AoII minimization and safety. 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 [24]. 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 AoI literature, as it can optimize the average age in multi-source systems [20]. In our case, it is equivalent to RR if $\varepsilon = 0$, and has the same issues in large networks with many



(a) AoII violation probability ($\Theta_{\max} = 0$).



(b) AoII violation probability ($\Theta_{\max} = 5$).

Fig. 6: AoII violation probability as a function of ρ , $N = 20$.

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

We consider the performance of the protocols as the AoII violation probability $V(\Theta_{\max})$, which corresponds to the fraction of time that the nodes spend with an AoII higher



(a) AoII violation probability ($\Theta_{\max} = 0$).



(b) AoII violation probability ($\Theta_{\max} = 5$).

Fig. 7: AoII violation as a function of N with $\rho = 0.5$.

than the threshold value $\Theta_{\max}$. 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, $\varepsilon = 0.05$, and ideal feedback, i.e., $\varepsilon_f = 0$.

Fig. 6 shows the violation probability as a function of the offered load ρ , considering a system with $N = 20$ nodes with the same activation probability and $\varepsilon = 0.05$. 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 extremely 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.

We can also consider the performance of the schemes as a function of the number of nodes N , considering a high constant load $\rho = 0.5$. As Fig. 7a shows, DELTA is remarkably robust to an increased number of nodes, and $V(0)$ remains approximately the same for up to $N = 50$ nodes, while scheduled algorithms gradually degrade due to the longer duration between subsequent transmission opportunities for the same node. On the other hand, Fig. 7a shows that $V(5)$ slightly increases even for DELTA, as collisions become harder to handle, but the performance gap between it and the benchmark protocol widens as the number of nodes increases. As before, DELTA is remarkably robust to different values of K : the optimistic and pessimistic models lead to approximately the




Fig. 8: $V(5)$ as a function of ν with $\rho = 0.5$.




Fig. 9: $V(5)$ as a function of ε_f with $\rho = 0.5$.

same performance, as does setting a fixed $K = \frac{5}{2}N$.

We then consider the robustness to errors in the estimated activation rates: we set a load $\rho = 0.5$, and randomly sampled 100 activation probability vectors $\lambda \sim \mathcal{U}(\frac{\rho}{N} - \nu, \frac{\rho}{N} + \nu)$. The input to DELTA was then the average vector, with growing differences among nodes as ν increased. The resulting AoII violation probability is shown in Fig. 8: all protocols are robust to this type of disruption, and in particular, DELTA is insensitive to changes in the activation probabilities in its BT phase, as long as the overall load is approximately correct.

Finally, Fig. 9 shows the robustness of DELTA to imperfect feedback: considering the adaptation of feedback messages, the protocol degrades gracefully, maintaining an advantage over scheduled mechanisms for $\varepsilon_f \leq 0.1$. A load $\rho = 0.5$ is already close to DELTA's saturation point, with collisions becoming a frequent occurrence, and our scheme comes out ahead even in this pessimistic performance evaluation: its robustness to feedback error holds in less intense traffic conditions, with a bigger performance improvement.

V. 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 most significant parameter settings. 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: formal verification of the protocol using DEL under more challenging settings, such as imperfect feedback or correlated activation, as well as extension to anomalies that are modeled as a more complex N -state Markov process.

REFERENCES

- [1] C. De Alwis, A. Kalla, Q.-V. Pham, P. Kumar, K. Dev, W.-J. Hwang, and M. Liyanage, "Survey on 6G frontiers: Trends, applications, requirements, technologies and future research," *IEEE Open Journal of the Communications Society*, vol. 2, pp. 836–886, Apr. 2021.
- [2] S. Kaul, M. Gruteser, V. Rai, and J. Kenney, "Minimizing age of information in vehicular networks," in *Proceedings of the 8th International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 2011, pp. 350–358.
- [3] A. Kosta, N. Pappas, and V. Angelakis, "Age of information: A new concept, metric, and tool," *Foundations and Trends in Networking*, vol. 12, no. 3, pp. 162–259, Nov. 2017.
- [4] A. Maatouk, S. Kriouile, M. Assaad, and A. Ephremides, "The age of incorrect information: A new performance metric for status updates," *IEEE/ACM Transactions on Networking*, vol. 28, no. 5, pp. 2215–2228, Jul. 2020.
- [5] Z. Lu, R. Li, K. Lu, X. Chen, E. Hossain, Z. Zhao, and H. Zhang, "Semantics-empowered communications: A tutorial-cum-survey," *IEEE Communications Surveys & Tutorials*, vol. 26, no. 1, pp. 41–79, Mar. 2024.
- [6] O. Ayan, S. Hirche, A. Ephremides, and W. Kellerer, "Optimal finite horizon scheduling of wireless networked control systems," *IEEE/ACM Transactions on Networking*, vol. 32, no. 2, pp. 927 – 942, Apr. 2024.
- [7] N. Abramson, "The ALOHA system: Another alternative for computer communications," in *Proceedings of the Fall Joint Computer Conference (AFIPS)*. ACM, Nov. 1970, pp. 281–285.
- [8] 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. 27 190–27 204, May 2018.
- [9] O. T. Yavascan and E. Uysal, "Analysis of slotted ALOHA with an age threshold," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1456–1470, May 2021.
- [10] A. Munari and E. Uysal, "Information freshness in random access channels for IoT systems," in *Proceedings of the International Balkan Conference on Communications and Networking (BalkanCom)*. IEEE, Sep. 2021, pp. 112–116.
- [11] L. Li, Y. Dong, C. Pan, and P. Fan, "Timeliness of wireless sensor networks with random multiple access," *Journal of Communications and Networks*, vol. 25, no. 3, pp. 405–418, Jun. 2023.
- [12] M. A. Abd-Elmagid, N. Pappas, and H. S. Dhillon, "On the role of age of information in the Internet of Things," *IEEE Communications Magazine*, vol. 57, no. 12, pp. 72–77, Dec. 2019.
- [13] J. Yang, D. B. Minturn, and F. T. Hady, "When poll is better than interrupt," in *Proceedings of the 10th Conference on File and Storage Technologies (FAST)*, vol. 12. USENIX, Feb. 2012.
- [14] H. Wang, D. Zhang, Z. Yang, and W. Li, "An adaptive scheduling framework for distributed key-value stores using RDMA," in *Proceedings of the 8th Annual International Conference on Network and Information Systems for Computers (ICNISC)*. IEEE, Sep. 2022, pp. 605–611.
- [15] L. Scheuven, T. Höfler, P. Schulz, N. Franchi, A. N. Barreto, and G. P. Fettweis, "State-aware resource allocation for wireless closed-loop control systems," *IEEE Transactions on Communications*, vol. 69, no. 10, pp. 6604–6619, Oct. 2021.
- [16] A. Maatouk, M. Assaad, and A. Ephremides, "On the age of information in a CSMA environment," *IEEE/ACM Transactions on Networking*, vol. 28, no. 2, pp. 818–831, Apr. 2020.
- [17] A. Baltag and L. S. Moss, "Logics for epistemic programs," *Synthese*, vol. 139, pp. 165–224, Mar. 2004.
- [18] H. Pan, T.-T. Chan, J. Li, and V. C. Leung, "Age of information with collision-resolution random access," *IEEE Transactions on Vehicular Technology*, vol. 71, no. 10, pp. 11 295–11 300, Oct. 2022.
- [19] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 5, pp. 1183–1210, May 2021.
- [20] A. M. Bedewy, Y. Sun, S. Kompella, and N. B. Shroff, "Optimal sampling and scheduling for timely status updates in multi-source networks," *IEEE Transactions on Information Theory*, vol. 67, no. 6, pp. 4019–4034, Jun. 2021.
- [21] J. Holm, F. Chiariotti, A. E. Kalør, B. Soret, T. B. Pedersen, and P. Popovski, "Goal-oriented scheduling in sensor networks with application timing awareness," *IEEE Transactions on Communications*, vol. 71, no. 8, pp. 4513–4527, Aug. 2023.
- [22] M. Moradian, A. Dadlani, A. Khonsari, and H. Tabassum, "Age-aware dynamic frame slotted aloha for machine-type communications," *IEEE Transactions on Communications*, vol. 72, no. 5, pp. 2639–2654, May 2024.
- [23] 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 *Proceedings of the Global Communications Conference (GLOBECOM)*. IEEE, Dec. 2023, pp. 3573–3578.
- [24] Z. Jiang, B. Krishnamachari, X. Zheng, S. Zhou, and Z. Niu, "Timely status update in wireless uplinks: Analytical solutions with asymptotic optimality," *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3885–3898, Apr. 2019.
- [25] J. Thota, N. F. Abdullah, A. Doufexi, and S. Armour, "V2V for vehicular safety applications," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 6, pp. 2571–2585, Jun. 2020.
- [26] 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, Nov. 2021.
- [27] Z. Ning, P. Dong, X. Wang, X. Hu, L. Guo, B. Hu, Y. Guo, T. Qiu, and R. Y. Kwok, "Mobile edge computing enabled 5G health monitoring for Internet of medical things: A decentralized game theoretic approach," *IEEE Journal on Selected Areas in Communications*, vol. 39, no. 2, pp. 463–478, Feb. 2021.
- [28] E. Uhlemann, "Connected-vehicles applications are emerging," *IEEE Vehicular Technology Magazine*, vol. 11, no. 1, pp. 25–96, Mar. 2016.
- [29] H. Hui, S. Wei, and W. Chen, "Fresh multiple access: A unified framework based on large models and mean-field approximations," *IEEE/KICS Journal of Communications and Networks*, vol. 25, no. 5, pp. 598–620, Oct. 2023.
- [30] 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 Journal on Selected Areas in Information Theory*, vol. 4, pp. 651–666, Nov. 2023.
- [31] A. Munari, "Monitoring IoT sources over random access channels: Age of incorrect information and missed detection probability," in *Proceedings of the International Conference on Communications (ICC)*. IEEE, Jun. 2024, pp. 207–213.
- [32] 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 *Proceedings of the International Conference on Communications (ICC)*. IEEE, May 2023, pp. 1688–1693.
- [33] F. Malandrino, C. F. Chiasserini, and G. Di Giacomo, "Efficient distributed DNNs in the mobile-edge-cloud continuum," *IEEE/ACM Transactions on Networking*, vol. 31, no. 4, pp. 1702–1716, Aug. 2023.
- [34] L. Huang, Y. Wu, J. M. Parra-Ullauri, R. Nejabati, and D. Simeonidou, "AI model placement for 6G networks under epistemic uncertainty estimation," *arXiv preprint arXiv:2402.11245*, Feb. 2024.
- [35] X. Chen and H. Deng, "Analysis of cryptographic protocol by dynamic epistemic logic," *IEEE Access*, vol. 7, pp. 29 981–29 988, Mar. 2019.
- [36] J. J. Kline, "Evaluations of epistemic components for resolving the muddy children puzzle," *Economic Theory*, vol. 53, no. 1, pp. 61–83, May 2013.
- [37] J. Plaza, "Logics of public communications," *Synthese*, vol. 158, pp. 165–179, Jul. 2007.