

Context-aware Status Updating: Wireless Scheduling for Maximizing Situational Awareness in Safety-critical Systems

Tasmeen Zaman Ornee Md Kamran Chowdhury Shisher

Dept. of ECE
Auburn University
Auburn, AL, USA
tzo0017@auburn.edu

Dept. of ECE
Auburn University
Auburn, AL, USA
mzs0153@auburn.edu

Clement Kam

Information Technology Division
U.S. Naval Research Laboratory
Washington, DC, USA
clement.kam@nrl.navy.mil

Yin Sun

Dept. of ECE
Auburn University
Auburn, AL, USA
yzs0078@auburn.edu

Abstract—In this study, we investigate a context-aware status updating system consisting of multiple sensor-estimator pairs. A centralized monitor pulls status updates from multiple sensors that are monitoring several safety-critical situations (e.g., carbon monoxide density in forest fire detection, machine safety in industrial automation, and road safety). Based on the received sensor updates, multiple estimators determine the current safety-critical situations.

Due to transmission errors and limited communication resources, the sensor updates may not be timely, resulting in the possibility of misunderstanding the current situation. In particular, if a dangerous situation is misinterpreted as safe, the safety risk is high. In this paper, we introduce a novel framework that quantifies the penalty due to the unawareness of a potentially dangerous situation. This situation-unaware penalty function depends on two key factors: the Age of Information (AoI) and the observed signal value. For optimal estimators, we provide an information-theoretic bound of the penalty function that evaluates the fundamental performance limit of the system. To minimize the penalty, we study a pull-based multi-sensor, multi-channel transmission scheduling problem. Our analysis reveals that for optimal estimators, it is always beneficial to keep the channels busy. Due to communication resource constraints, the scheduling problem can be modelled as a Restless Multiarmed Bandit (RMAB) problem. By utilizing relaxation and Lagrangian decomposition of the RMAB, we provide a lowcomplexity scheduling algorithm which is asymptotically optimal. Our results hold for both reliable and unreliable channels. Numerical evidence shows that our scheduling policy can achieve up to 100 times performance gain over periodic updating and up to 10 times over randomized policy.

Index Terms—safety, age of information, Markov decision process, estimation

I. INTRODUCTION

A broad range of safety-critical systems is ubiquitous across the world. For instance, in industrial automation, it is essential to continuously monitor the safety of various machines [1]. In patient health monitoring, precise tracking of the glucose level or the heart rate is imperative to swiftly implement precautionary measures when they are required [2]. In disaster monitoring, it is important to promptly monitor any consistent changes in temperature or humidity, as they could indicate a possible disaster [3]. In these safety-critical situations, the monitoring system needs timely access and accurately interpret

This work was supported in part by NSF grant CNS-2239677, ARO grant W911NF-21-1-0244, and ONR.

979-8-3503-2181-4/23/\$31.00 ©2023 IEEE

the states of remote systems. Any misunderstanding of the system state can lead to severe consequences.

In practice, multiple sensors are required to track various safety-critical situations. One challenge to continuously monitor these sensor measurements in real-time is the limited capacity of the communication medium. Moreover, some sensors may have more crucial content than others and hence need more attention. In this context, we adopt a pull-based system where a centralized monitor selects sensors and requests information when required. This selective retrieval of information ensures that the system receives essential information promptly while minimizing unnecessary resource consumption.

In this paper, we consider a discrete-time pull-based status updating system consisting of multiple sensors monitoring the status of different safety-critical situations. At every time slot, the selected sensors transmit their updates to a receiver through multiple unreliable channels. In the receiver, multiple estimators utilize the sensor updates to determine the current status of the safety-critical situations. Due to transmission errors, the sensor updates may not be fresh. One performance metric that characterizes data freshness is the age of information (AoI) [4]. Let $U(t)$ be the generation time of the freshest observation delivered to the receiver by time t . The AoI, as a function of t , is defined as $\Delta(t) = t - U(t)$ which exhibits a linear growth with time t and drops down to a smaller value whenever a fresher observation is delivered. In many real-time applications, it is important to consider AoI for making the scheduling decision. However, the time difference represented by AoI can only capture the timeliness of the information but it cannot capture its significance. This is particularly relevant in safety-critical situations where misunderstanding about the situation can lead to significant

performance loss. Hence, relying solely on Aol-based decision-making is not perfect. If we consider signal observation along with Aol in decision-making, then the incurred performance loss can be significantly improved. One key observation in this study is that any misinterpretation of a dangerous situation yields a higher loss compared to the misinterpretation of a safe situation. Based on the above-mentioned insights, we introduce a framework for quantifying the *cost of a dangerous situation* that characterizes the performance loss caused by situational unawareness.

The goal of this paper is to find the optimal scheduling policy to select sensors and to request observations while improving the system performance. The contributions of this paper are as follows:

- We introduce a novel framework for estimating the current status of a safety-critical system. In this framework, we adopt a general loss function $L(y, \hat{y})$ that quantifies the incurred loss in wrongly estimating the actual safety level y as $\hat{y}$. The loss $L(\text{dangerous}, \text{safe})$ is higher than $L(\text{safe}, \text{dangerous})$. This distinction can not be captured by the traditional loss functions such as 0-1 loss, quadratic loss, and logarithmic loss. By adopting appropriate loss functions L , our framework can be applied to health, safety, and security monitoring.
- To measure the performance of safety-critical systems, we propose a penalty function that represents the expected loss L given the Aol and the latest observation (see Section III-C). We also provide an informationtheoretic lower bound of the penalty function by using L -conditional entropy [5]–[7]. This bound represents the fundamental performance limit of a safety-critical system. The entropy-based freshness metric in our study can significantly contribute to real-time applications such as estimation, inference, and perception. Earlier metrics such as Age of Incorrect Information (AolI) [8], Age of Synchronization (AoS) [9], Urgency of Information (Uol) [10], Version Aol [11], Aol at Query (QAol) [12], Value of Information (Vol) [13], and Uncertainty of Information (Uol) [14] did not quantify the fundamental performance limit of real-time applications. Moreover, most of the prior studies [8]–[12] exhibit a monotonic relationship with Aol, whereas some recent studies show that the performance of real-time applications may degrade nonmonotonically with Aol [7], [14], [15]. Our penalty function also allows the non-monotonic behavior with respect to Aol.
- We consider a multi-sensor, multi-channel pull-based status updating problem. Our findings demonstrate that when utilizing one-time slot transmission time and optimal estimators, it is always beneficial to keep the channels busy (see Theorem 1). However, channel resource limitations prevent all sensors from transmitting information continuously. To address this issue, we formulate the multi-sensor, multi-channel transmission

scheduling problem as a Restless Multi-armed Bandit (RMAB). We utilize relaxation and Lagrangian method to decompose the original problem into multiple separated Markov Decision Processes (MDPs). We solve each MDP by dynamic programming [16]. By utilizing the solution to the MDPs, we provide a low-complexity scheduling policy which is asymptotically optimal. Our results work for both reliable and unreliable channels.

- Numerical results illustrate that our multi-sensor, multichannel scheduling policy achieves up to 100 times performance gain over periodic updating policy and up to 10 times over randomized policy which randomly select sensors depending on the number of available channels.

II. RELATED WORK

There exists a large number of studies on minimizing linear and nonlinear Aol functions [17]–[22]. One limitation of Aol is that it only captures the timeliness of the information while neglecting the actual influence of the conveyed information. In order to address this, several performance metrics were introduced in conjunction with Aol [8]–[12], [14], [23]. Age of Incorrect Information (AolI) was introduced in [8] that is represented by a function of the age and the estimation error. In [9], Age of Synchronization (AoS) was considered along with Aol to measure the freshness of a local cache. In [10], the authors proposed Urgency of Information (Uol) that captures the context-dependence of the status information along with Aol. Version Aol was introduced in [11] that represents how many versions are out-of-date at the receiver, compared to the transmitter. An Aol at Query (QAol) metric was introduced in [12] to capture the freshness only when required in a pull-based communication system. In addition, several research papers studied information-theoretic measures to evaluate the impact of information freshness along with information content [13], [14], [17], [24], [25]. In [13], [17], [24], [25], the authors employed Shannon's mutual information to quantify the information carried by received data messages regarding the current signal value at the source and used Shannon's conditional entropy to measure the uncertainty about the current signal value. Based on the studies of [13], [17], [24], [25], the authors in [14] termed Uncertainty of Information (Uol) by using the Shannon's entropy. However, there exists a disparity between these information-theoretic metrics and the performance of real-time applications such as remote estimation and inference. In [7], [15], [26], a generalized conditional entropy associated with a loss function L , or L -conditional entropy $H_L(Y_t|X_{t-\Delta(t)})$ was utilized to address this disparity, where Y_t is the true state of the source and $X_{t-\Delta(t)}$ is the observed value. Building upon the insights of [7], [15], [26], we utilized L -conditional entropy $H_L(Y_t|X_{t-\Delta(t)} = x, \Delta(t) = \delta)$ given both the Aol δ and the observed value x to measure the impact of the Aol and the information content in remote estimation and inference.

In addition, there exists numerous papers on Aol-based sampling and scheduling [7], [17], [19], [26]–[32]. In [17], sampling policies for optimizing non-linear Aol functions were studied. A joint sampling and scheduling problem to minimize

monotonic Aol functions was considered in [19]. A Whittle index-based scheduling algorithm to minimize Aol for stochastic arrivals was considered in [28]. In [14], the authors proposed a Whittle index-based scheduling policy to minimize the Uol modeled as Shanon entropy. Optimal scheduling policies for both single and multi-source systems were studied and a Whittle index policy was proposed for multi-source cases in [7]. A Whittle index policy for both signal-aware and signalagnostic scheduling was reported in [29]. A remote estimation system with multiple IoT sensors monitoring multiple Wiener processes was studied and a Max- k policy was proposed in [33]. Besides Whittle index-based policies that require an indexability condition, non-indexable scheduling policies were also studied in [26], [30]–[32], [34]. In this paper, because of the complicated nature of state transition along with erasure channels, we do not provide indexability. However, we provide




Fig. 1: A multi-sensor, multi-channel vehicle safety monitoring system.

a “Net-gain Maximization Policy” developed in [26], [34]. In addition, by utilizing information-theoretic approach, we show that it is always beneficial to keep the channels busy for optimal estimators. Our scheduling policy is designed for pullbased communication model where the scheduling decision is based on Aol and observed signal and the developed policy is asymptotically optimal.

III. MODEL, METRIC, AND FORMULATION

A. System Model

Let us consider the status updating system depicted in Figure 1, where N sensors transmit crucial status updates through M unreliable wireless channels to a common receiver. Each sensor n monitors a Markov signal $X_{n,t}$ representing the status of a safety-critical system. For instance, $X_{n,t}$ might represent the position of a vehicle on the road, or the joint angles of a robotic arm within a factory environment. A hazardous situation arises when the vehicle veers off the road or the robotic arm approaches a nearby object. We use $Y_{n,t}$ to quantify the level of danger for the safety-critical system, which is a function of the system status $X_{n,t}$. In practice, $Y_{n,t}$ can

be used to represent whether the vehicle is encroaching upon the road shoulder or the spatial distance between the robotic arm and the object. The receiver estimates the danger level signals $Y_{n,t}$ to ensure awareness of the hazards in the safety-critical systems.

We consider a pull-based updating mechanism [35] where the receiver requests status updates from the sensors whenever it is unsure about the situation. In response to the pull request, each sensor n generates and submits a time-stamped updating message $(X_{n,t}, t)$ to one wireless channel. We assume that it takes one-time slot for the transmission of a message update to the receiver. Due to wireless channel fading, the transmission of the status updates becomes unreliable. Let p_n be the probability of a successful transmission from sensor n , irrespective of the selected wireless channel.

Due to transmission errors, the information received by the receiver will be stale and is represented by $X_{n,t-\Delta_n(t)}$ that is generated $\Delta_n(t)$ times ago. The time different $\Delta_n(t)$ is usually called *age of information (Aol)* [4], which represents the staleness of the status of the n safety-critical system available at the receiver. At each time slot t , the Aol evolution of the n -th system is given by

$$\Delta_n(t+1) = \Delta_n(t) + 1, \text{ with probability } 1 - p_n, \quad (1) \quad \Delta_n(t+1) = 0, \text{ with probability } p_n.$$

B. Loss Model for Situational Awareness

Based on the latest available information, the n -th estimator outputs $a = \phi_n(\Delta_n(t), X_{n,t-\Delta_n(t)}) \in A$, where $\phi : N \times X \rightarrow A$ is a function of Aol $\Delta_n(t) \in N$ and the received

observation $X_{n,t-\Delta_n(t)} \in X$. The danger associated with the safety-critical system is characterized by a loss function $L : Y \times A \rightarrow R$, where $L(y, a)$ is the incurred loss if $Y_{n,t} = y$ is the actual safety level of the surrounding environment and a is estimated output of the safety level. Essentially, $L(\cdot, \cdot)$ serves as a metric to assess the *cost of potential danger* within a safety-critical system. To better understand the behavior of the loss L , we provide the following example that illustrate the impact of wrong estimation of the system state on L .

Example Consider a road safety monitoring system that detects instances when any car veers off the road. The sensors need to operate with high sensitivity to accurately communicate the measured variables in real-time. Let $Y_{n,t} = \{\text{danger}, \text{safe}\}$ denote the safety measure of a car based on its position, speed, etc. If $y = \text{danger}$ and $a = \text{safe}$, then the loss $L(\text{danger}, \text{safe})$ would be significantly high. This is because if the car is not within the safe region and the monitoring system wrongly estimates it, there could be a serious damage. However, if $y = \text{safe}$ but $a = \text{danger}$, then the loss $L(\text{safe},$

danger) would have less impact. This is because even though the estimation is wrong, the car is still within the safe region, hence, it does not impact much.

The well-known loss functions such as 0-1 loss, quadratic loss, and logarithmic loss cannot address safety issues based on situational awareness within critical systems. The loss function L in (2) is more general than the existing loss functions. By designing this unified loss function L that incorporates the knowledge of the surrounding situation along with age, we can effectively capture and tackle safety-critical issues.

C. Information-theoretic Metric for Situational Awareness


The performance of the safety-critical system for sensor n at time slot t is defined by the expected loss for a given state $(\Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x)$ which can be expressed as the following penalty function:

$$q_n(\delta, x) = E[L(Y_{n,t}, \phi_n(\Delta_n(t), X_{n,t-\Delta_n(t)})) |$$


$$\Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x], \quad (2) \text{ where } \phi_n(\cdot, \cdot) \text{ is}$$

any function that maps from $N \times X$ to A . Now, consider the following optimization problem:

$$\min_{\phi_n} E[L(Y_{n,t}, \phi_n(\Delta_n(t), X_{n,t-\Delta_n(t)})) | \Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x] \quad (3)$$



(a) Gridworld



(b) Penalty vs Aol

Fig. 2: (a) Gridworld environment and (b) Penalty $(q_n(\delta, x))$ vs Aol (δ) for four different given observation.

Let $\hat{\phi}_n$ be the optimal estimator that solves the optimal estimation problem in (3). By substituting this optimal estimator into (2), we get $q_n(\delta, x) = E[L(Y_{n,t}, \hat{\phi}_n(\Delta_n(t), X_{n,t-\Delta_n(t)})) |$

$$\Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x], \quad (4) \text{ which is a}$$

lower bound of $E[L(Y_{n,t}, a) | \Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x]$ for any $a \in A$. This penalty function in (4) is closely related to the concept of generalized entropy [5], [6] or specifically, the L -entropy [7] of a random variable $Y_{n,t}$ given by

$$H_L(Y_{n,t}) = \min_{a \in A} E[L(Y_{n,t}, a)]. \quad (5)$$

Furthermore, L -conditional entropy of $Y_{n,t}$ given $\Delta_n(t) = \delta$ and $X_{n,t-\Delta_n(t)} = x$ can be defined as [5]–[7]

$$\begin{aligned} H_L(Y_{n,t} | \Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x) &= \min_{a \in A} E[L(Y_{n,t}, a) | \Delta_n(t) \\ &= \delta, X_{n,t-\Delta_n(t)} = x], \\ &= E[L(Y_{n,t}, \hat{\phi}_n(\Delta_n(t), X_{n,t-\Delta_n(t)})) | \Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x]. \end{aligned} \quad (6)$$

From (4) and (6), it is evident that $q_n(\delta, x) = H_L(Y_{n,t} | \Delta_n(t) = \delta, X_{n,t-\Delta_n(t)} = x)$. (7)

For the optimal estimator $\hat{\phi}_n$ is indeed L -conditional entropy which is an information-theoretic lower bound of $q_n(\delta, x)$. It represents the fundamental performance limit that characterizes the performance degradation due to the lack of the knowledge of the situation. The proposed metrics in prior works, i.e., AoI, Vol, AoS, QAol cannot explain this information theoretic bound. In addition, for any general estimator output $a \in A$, $q_n(\delta, x)$ can be represented as the L -conditional cross-entropy. Due to space limitation, the details are relegated to our future submission.

D. Non-monotonic Information Aging

Our analysis reveals that $q_n(\delta, x)$ can be a non-monotonic function of the age, particularly when the knowledge of the surrounding situation is taken into consideration which is illustrated in Figure 2(b). To do this experiment, we consider a safety-critical system where N robots are moving in a

gridworld with 5 rows and 12 columns, demonstrated in Figure 2(a). The observed state $X_{n,t}$ of robot n is represented by two variables: the position $S_{n,t}$ of robot n at time t and its moving direction $a_{n,t}$ at time t and $Y_{n,t} = \{safe, cautious, dangerous\}$ denotes the safety level. In the gridworld in Figure 2(a), the states $X_{n,t}$ in row 1, 2, and 3 are *safe*, row 4 states are *cautious*, and row 5 states are *dangerous*. Row 3 is close to the boundary region between *safe* and *cautious*. Let (x, y) denotes the position of a robot where x is the row and y is the column. The available moving directions $a_{n,t}$ for row 2, 3, and 4 are *up*, *down*, *left*, and *right*. For row 1, the *up* is not available and for row 5, the *down* is not available. If robot n is in the leftmost position, then $a_{n,t} = \text{left}$ means it will stay in the same position, similar criteria is applied for the rightmost position. The probability of moving from one row to the adjacent row is 0.05 (*up* or *down*) and the probability of staying in the same row is 0.95 (*left* or *right*). The losses considered in Figure 2(b) are:

$L(\text{cautious}, \text{safe}) = 50, L(\text{safe}, \text{cautious}) = 10, L(\text{dangerous}, \text{safe}) = 200, L(\text{safe}, \text{dangerous}) = 10, L(\text{dangerous}, \text{cautious}) = 50, L(\text{cautious}, \text{dangerous}) = 20$, and $L(\text{dangerous}, \text{dangerous}) = L(\text{cautious}, \text{cautious}) = L(\text{safe}, \text{safe}) = 0$. We consider optimal estimator of (3) in this experiment.

From Figure 2(b), we observe that when a robot is in a *safe* region and far from the *safe* and *cautious* boundary which is represented by the curve for given $X_{n,t} = (1,3)$, *right*, the penalty is initially close to zero for small AoI values and increases gradually with increasing age. This phenomenon tells us that we do not need to update frequently when a robot is far from the boundary region. However, if the robot moves closer to the boundary between *safe* and *cautious* that is represented by the curve for given $X_{n,t} = (2,3)$, *down*, the penalty increases very quickly because of the uncertainty of its position in the subsequent time slots. With the increase in age, this curve approaches to its stationary distribution. In similar way, the other curves can be explained. This penalty curves are not necessarily monotonic with age. Hence, only considering the non-decreasing functions of the age is not sufficient for performance analysis of safety-critical systems. The proposed metrics in prior works, i.e, AoI, Vol, AoS, QAoI cannot explain this non-monotonicity with age.

E. Scheduling Policy and Problem Formulation

Let the scheduling policy is denoted by π_n where $\pi_n = (\mu_n(0), \mu_n(1), \dots)$ determines whether an observation is requested from sensor n at every time slot $t \in \mathbb{N}$. Let Π denotes the set of all causal scheduling policies in which every decision is made by using the current and history information available at the receiver. Because our system consists of M channels, π_n is required to hold for all t .

We aim to find an optimal scheduling policy that minimizes the time-average sum of expected penalty of the N sources over an infinite time-horizon T , which is formulated as

$$q_{\text{opt}} = \min_{\pi_n} \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \sum_{n=1}^N \mathbb{E}[\mu_n(t) L(X_{n,t})] \quad (9)$$

s.t. $\mu_n(t) \in \{0,1\}, t = 0,1,\dots$, where $L(X_{n,t})$ is the penalty incurred by source n at time t which is defined in (2), and q_{opt} is the optimum value of (8).

IV. PENALTY-MINIMIZATION: AN INFORMATION-THEORETIC VIEW

In Section III-C, we demonstrate that the penalty function $q_n(\delta, x)$ can be represented as L -conditional entropy. Leveraging this insight, we obtain that for optimal estimators, always sending updates benefits the system by reducing its average penalty of the system. To prove this result, we present the following useful lemma which illustrates that more information reduces the L -conditional entropy.

Lemma 1. For random variables X, Y , and Z , it holds that $H_L(Y|Z=z) \geq H_L(Y|X,Z=z)$, where $H_L(Y|Z=z) = \min_{a \in A} \mathbb{E}[L(Y,a)|Z=z]$, (10)

$a \in A$

$$H_L(Y|X,Z=z) = \sum_{x \in X} P(X=x|Z=z) H_L(Y|X=x,Z=z).$$

(11) Then we have the following theorem.

Theorem 1. If the packet transmission times are one-time slot, then for optimal estimators it is always better to keep the channels busy.

Due to space limitation, the proofs of Lemma 1 and Theorem 1 are relegated to our future submission.

Because problem (8)-(9) has a channel resource constraint, all of the sensors cannot submit their updates at every time slot when $N > M$. Therefore, we have to design an efficient scheduling policy that minimizes the time-average sum of the expected penalty of the N sources ensuring that constraint (9) is satisfied. We provide the details in the next section.

V. RESTLESS MULTI-ARMED BANDIT FORMULATION

Problem (8)-(9) is an RMAB problem where each source n is an arm and $(\Delta_n(t), X_{n,t-\Delta_n(t)})$ is the state of each arm n . To find an optimal solution to the RMAB problem is PSPACE hard [36]. A Whittle index policy is known to be asymptotically optimal for many RMAB problems [37]. However, it needs to satisfy a complicated condition called indexability. Due to the complicated nature of the state transitions and non-monotonic age-penalty functions along with erasure channels, it is difficult to establish indexability for our Algorithm 1 Net-gain Maximization Policy

- 1: At time $t = 0$:
- 2: Input λ^* which is the optimal solution to (14).
- 3: Input $\alpha_{n,\lambda^*}(\delta, x)$ in (19) for every source n .
- 4: For all time $t = 0, 1, \dots$,
- 5: Update $(\Delta_n(t), X_{n,t-\Delta_n(t)})$ for all source n .
- 6: Update current "gain" $\alpha_{n,\lambda^*}(\Delta_n(t), X_{n,t-\Delta_n(t)})$ for all source n .
- 7: Choose at most M sensors with highest positive "gain".

problem. Therefore, in this work, we provide a low-complexity algorithm that does not need to satisfy indexability. Next, we demonstrate that the developed policy is asymptotically optimal.

A. Relaxation and Lagrangian Decomposition

Following the standard relaxation and Lagrangian decomposition procedure for RMAB [38], the original problem in (8)-(9) is relaxed as

$$\begin{aligned} & \min_{\pi} q_{\text{opt}} \\ & \text{s.t.} \end{aligned} \quad (13)$$

where the relaxed constraint (13) only needs to be satisfied on average, whereas (9) is required to hold at any time t . To solve the relaxed problem (12)-(13), we take a dual cost $\lambda \geq 0$ (also known as Lagrange multiplier) for the relaxed constraint. The dual problem is given by

$$\sup_{\lambda \geq 0} q^-(\lambda), \quad (14)$$

where

$$q^-(\lambda) = \min_{\pi} \left[\sum_{n=1}^N \left(\mathbb{E}_{\pi} [q_n(\delta, x)] - \lambda \right) \right] \quad (15)$$

The term $\mathbb{E}_{\pi} [q_n(\delta, x)]$ in (15) does not depend on policy π and hence can be removed. For a given λ , problem (15) can be decomposed into N separated sub-problems and each sub-problem associated with source n is formulated as

$$\min_{\pi_n} \left[\mathbb{E}_{\pi_n} [q_n(\delta, x)] - \lambda \right] \quad (16)$$

where $q_n^-(\lambda)$ is the optimum value of (16), $\pi_n = (\mu_n(1), \mu_n(2), \dots)$ denotes a sub-scheduling policy for source n , and Π_n is the set of all causal sub-scheduling policies of source n .

VI. OPTIMAL POLICY VIA DYNAMIC PROGRAMMING

Given transmission cost λ , the per-arm problem (16) is an average-cost infinite horizon MDP with state $(\Delta_n(t), X_{n,t-\Delta_n(t)})$. We solve (16) by using dynamic programming [16]. The Bellman optimality equation for the MDP in (16) is

$$q_n^-(\lambda) = \min_{\mu} \left[\mathbb{E}_{\mu} [q_n(\delta, x)] - \lambda \right] \quad (17)$$

where $h_{n,\lambda}(\delta, x)$ is the relative-value function of the average-cost MDP and $Q_{n,\lambda}(\delta, x, \mu)$ is the relative action-value function defined as

$$Q_{n,\lambda}(\delta, x, \mu) = \mathbb{E}_{\mu} [q_n(\delta, x) - q_n^-(\lambda) + h_{n,\lambda}(\delta + 1, x), \text{ if } \mu = 0, q_n(\delta, x) - q_n^-(\lambda) +$$

$$(1 - p_n)h_{n,\lambda}(\delta + 1, x) \quad (18)$$

$$+ p_n \mathbb{E} [h_{n,\lambda}(1, X_{n,t-1}) | X_{n,t-\delta} = x] + \lambda, \text{ otherwise.}$$

The relative-value function $h_{n,\lambda}(\delta, x)$ can be computed by using relative value iteration algorithm for average-cost MDP [16]. Following [26], [34], define the "gain" $\alpha_{n,\lambda}(\delta, x)$ for choosing the action $\mu_{n,\lambda}(t)$ as $\alpha_{n,\lambda}(\delta, x) = Q_{n,\lambda}(\delta, x, 0) - Q_{n,\lambda}(\delta, x, 1)$. (19) Substituting (18) into (19), we get

$$\alpha_{n,\lambda}(\delta, x) = \mathbb{E} [h_{n,\lambda}(1, X_{n,t-1}) | X_{n,t-\delta} = x] - \mathbb{E} [h_{n,\lambda}(\delta + 1, x)] \quad (20)$$

By utilizing the "gain" $\alpha_{n,\lambda}(\delta, x)$ in (20), we obtain the optimal decision to the relaxed problem (12)-(13) at time t for every sensor n as $\mu_n(t) = \arg \max_{\mu} \alpha_{n,\lambda}(t) (\Delta_n(t), X_{n,t-\Delta_n(t)}, \mu)$, (21) $\mu \in \{0, 1\}$

where the dual cost is iteratively updated using the dual subgradient ascent method with step size $\beta > 0$ [39]:

$$\lambda(t) = \lambda(t-1) + \beta (q^-(\lambda(t-1)) - \lambda(t-1)) \quad (22)$$

Let λ^* be the optimal dual cost to problem (14) to which $\lambda(t)$ converges. We provide a low-complexity algorithm for solving problem (8)-(9) in Algorithm 1. We utilize the "gain" defined in (19) as the priority measurement for choosing action μ . Algorithm 1 takes optimal dual cost λ^* and the precomputed gain $\alpha_{n,\lambda^*}(\delta, x)$ associated with λ^* as input. Then, for all $t \geq 0$, the state $(\Delta_n(t), X_{n,t-\Delta_n(t)})$ and the associated "gain" $\alpha_{n,\lambda^*}(\Delta_n(t), X_{n,t-\Delta_n(t)})$ are updated. Finally, Algorithm 1 maximizes the "Net-gain" (total gain of all sensors) of the system at time t . This is done by selecting at most M sensors having the highest positive "gain" at time t . The "Net-gain Maximization Policy" in Algorithm 1 does not need to satisfy the indexability condition.

VII. ASYMPTOTIC OPTIMALITY

In this section, we demonstrate that the "Net-gain Maximization Policy" in Algorithm 1 is asymptotically optimal in the same asymptotic regime as the Whittle index policy [38].





Fig. 3: Normalized average penalty vs Number of sources (N) where Number of channels are $M = 10$ with success probability 0.95.

In this scenario, all N arms are generalized to N classes, and the number of arms in each class and the number of channels M are scaled by γ , while maintaining a constant ratio between them.

Let $Z_n^\gamma(\pi^*, \{\delta, x\}, \mu, t)$ be a fluid-scaling process with parameter γ that represents the expected number of class- n arms at state (δ, x) that takes action μ at time slot t under policy π^* . Consider the following expected long-term average cost



The policy π^* will be asymptotically optimal if $\lim_{\gamma \rightarrow \infty} Z_n^\gamma(\pi^*; t)/\gamma = 0$ for all $\pi \in \Pi$. In this sequel, we introduce the following *global attractor* [40].

Definition 1. Global attractor. An equilibrium point $Z_n^{\gamma^*}/\gamma$ under policy π^* is a global attractor for the process

$Z_n^\gamma(\pi^*; t)/\gamma$, if, for any initial point $Z_n^\gamma(\pi^*; 0)/\gamma$, the process $Z_n^\gamma(\pi^*; t)/\gamma$ converges to $Z_n^{\gamma^*}/\gamma$.

Theorem 2. Under Definition 1, the policy π^* is asymptotically optimal. Therefore, $\lim_{\gamma \rightarrow \infty} Z_n^\gamma(\pi^*; t)/\gamma = 0$, where π^{opt} is the optimal policy for the original RMAB problem (8)-(9).

Due to space limitation, the proof of Theorem 2 is relegated to our future submission.

VIII. NUMERICAL RESULTS

In this section, we evaluate the performance of the following policies:

- Periodic Updating: The sensors generate updates at every time slot and store in a FIFO queue. Whenever a channel is available, an update from the queue is sent.
- Randomized Policy: If M channel resources are available, this policy randomly selects at most M sensors.
- Net-gain Maximization Policy: See Algorithm 1.

We consider the same experimental setup of Figure 2 where 5 robots follow a deterministic policy (they follow a fixed path). The cost associated with these 5 robots is zero at every time slot because given an initial state, the position of these robots can be uniquely determined by following the deterministic policy. The goal of the other $N - 5$ robots is to move and scan the environment (e.g., Mars Rovers [41]) and send updates when requested. We do not consider any termination state for these robots, the goal is to keep scanning for infinite-time horizon. Our system consists of $M = 10$ erasure channels and the success probability is 0.95.

The performance comparison of the three policies mentioned above is provided in Figure 3. The normalized average penalty in Figure 3 is obtained by dividing time-average cost by the number of robots. From the figure, until $N \leq M$, all of the three policies show the same performance. Whenever $N > M$, periodic updating starts getting worse because the queue length is getting higher. In our simulation, we have used a buffer size of 20 for periodic updating. Moreover, the randomized policy randomly selects at most 5 sensors for sending updates, whereas the net-gain maximization policy makes the decision in a smarter way by considering the Aol and the state of the surrounding situation. The performance gain of the net-gain maximization policy is up to 100 times compared to periodic updating and up to 10 times compared to the randomized policy.

IX. CONCLUSION

We address the importance of situational awareness in safety-critical systems. The general loss function L have practical importance and appropriate design of L can address many safety-critical issues. In future we will study systems where multiple sensors can arrive and leave the system at any time. Another interesting direction is to consider a finite time horizon problem where there is a termination state while encountering a danger.

REFERENCES

- [1] A. Grau, M. Indri, L. L. Bello, and T. Sauter, "Industrial robotics in factory automation: From the early stage to the internet of things," in *IEEE IECON*, 2017, pp. 6159–6164.
- [2] S. Abdulmalek, A. Nasir, W. A. Jabbar, M. A. Almuhaaya, A. K. Bairagi, M. A.-M. Khan, and S.-H. Kee, "IoT-based healthcare-monitoring system towards improving quality of life: A review," in *Healthcare*, vol. 10, no. 10, 2022, p. 1993.
- [3] M. Seenivasan, M. Arularasu, K. Senthilkumar, and R. Thirumalai, "Disaster prevention and control management in automation: a key role in safety engineering," *Procedia Earth and Planetary Science*, vol. 11, pp. 557–565, 2015.
- [4] S. Kaul, R. D. Yates, and M. Gruteser, "Real-time status: How often should one update?" in *IEEE INFOCOM*, 2012.
- [5] A. P. Dawid, "Coherent measures of discrepancy, uncertainty and dependence, with applications to bayesian predictive experimental design," *Department of Statistical Science, University College London*, vol. 139, 1998.
- [6] F. Farnia and D. Tse, "A minimax approach to supervised learning," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [7] M. K. C. Shisher and Y. Sun, "How does data freshness affect real-time supervised learning?" in *ACM MobiHoc*, 2022, pp. 31–40.
- [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, p. 2215–2228, oct 2020.
- [9] J. Zhong, R. D. Yates, and E. Soljanin, "Two freshness metrics for local cache refresh," in *IEEE ISIT*, 2018, pp. 1924–1928.
- [10] X. Zheng, S. Zhou, and Z. Niu, "Urgency of information for contextaware timely status updates in remote control systems," *IEEE Trans. Wirel. Commun.*, vol. 19, no. 11, pp. 7237–7250, 2020.
- [11] R. D. Yates, "The age of gossip in networks," in *IEEE ISIT*, 2021, pp. 2984–2989.
- [12] J. Holm, A. E. Kalør, F. Chiariotti, B. Soret, S. K. Jensen, T. B. Pedersen, and P. Popovski, "Freshness on demand: Optimizing age of information for the query process," in *IEEE ICC*, 2021, pp. 1–6.
- [13] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, "Age and value of information: Non-linear age case," in *IEEE ISIT*, 2017, pp. 326–330.
- [14] G. Chen, S. C. Liew, and Y. Shao, "Uncertainty-of-information scheduling: A restless multi-armed bandit framework," *IEEE Trans. Inf. Theory*, 2022.

- [15] M. K. C. Shisher, H. Qin, L. Yang, F. Yan, and Y. Sun, "The age of correlated features in supervised learning based forecasting," in *IEEE INFOCOM Workshops*, 2021, pp. 1–8.
- [16] D. P. Bertsekas *et al.*, "Dynamic programming and optimal control 3rd edition, volume ii," *Belmont, MA: Athena Scientific*, vol. 1, 2011.
- [17] Y. Sun and B. Cyr, "Sampling for data freshness optimization: Nonlinear age functions," *J. Commun. Netw.*, vol. 21, pp. 204–219, 2019.
- [18] Y. Sun, E. Uysal-Biyikoglu, R. D. Yates, C. E. Koksall, and N. B. Shroff, "Update or wait: How to keep your data fresh," *IEEE Trans. Inf. Theory*, vol. 63, no. 11, pp. 7492–7508, 2017.
- [19] 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.
- [20] T. Z. Ornee and Y. Sun, "Sampling and remote estimation for the ornstein-uhlenbeck process through queues: Age of information and beyond," *IEEE/ACM Trans. Netw.*, vol. 29, no. 5, p. 1962–1975, oct 2021.
- [21] Y. Sun, Y. Polyanskiy, and E. Uysal, "Sampling of the wiener process for remote estimation over a channel with random delay," *IEEE Trans. Inf. Theory*, vol. 66, no. 2, pp. 1118–1135, 2020.
- [22] R. D. Yates, Y. Sun, D. R. Brown, S. K. Kaul, E. Modiano, and S. Ulukus, "Age of information: An introduction and survey," *IEEE J. Sel. Areas Commun.*, vol. 39, no. 5, pp. 1183–1210, 2021.
- [23] N. Pappas and M. Kountouris, "Goal-oriented communication for realtime tracking in autonomous systems," in *IEEE ICAS*, 2021, pp. 1–5.
- [24] Z. Wang, M.-A. Badiu, and J. P. Coon, "A framework for characterizing the value of information in hidden markov models," *IEEE Trans. Inf. Theory*, vol. 68, no. 8, pp. 5203–5216, 2022.
- [25] T. Soleymani, S. Hirche, and J. S. Baras, "Optimal self-driven sampling for estimation based on value of information," in *IEEE WODES*, 2016, pp. 183–188.
- [26] M. K. C. Shisher, B. Ji, I. Hou, Y. Sun *et al.*, "Learning and communications co-design for remote inference systems: Feature length selection and transmission scheduling," *arXiv preprint arXiv:2308.10094*, 2023.
- [27] I. Kadota, A. Sinha, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, "Scheduling policies for minimizing age of information in broadcast wireless networks," *IEEE/ACM Trans. Netw.*, vol. 26, no. 6, pp. 2637–2650, 2018.
- [28] Y.-P. Hsu, "Age of information: Whittle index for scheduling stochastic arrivals," in *IEEE ISIT*, 2018, pp. 2634–2638.
- [29] T. Z. Ornee and Y. Sun, "A Whittle index policy for the remote estimation of multiple continuous Gauss-Markov processes over parallel channels," accepted by *ACM MobiHoc* 2023.
- [30] G. Xiong, X. Qin, B. Li, R. Singh, and J. Li, "Index-aware reinforcement learning for adaptive video streaming at the wireless edge," in *ACM MobiHoc*, 2022, pp. 81–90.
- [31] Y. Zou, K. T. Kim, X. Lin, and M. Chiang, "Minimizing age-of-information in heterogeneous multi-channel systems: A new partialindex approach," in *ACM MobiHoc*, 2021, pp. 11–20.
- [32] Y. Chen and A. Ephremides, "Scheduling to minimize age of incorrect information with imperfect channel state information," *Entropy*, vol. 23, no. 12, p. 1572, 2021.
- [33] J. Yun, A. Eryilmaz, J. Moon, and C. Joo, "Remote estimation for dynamic iot sources under sublinear communication costs," *IEEE/ACM Trans. Netw.*, 2023.
- [34] G. Chen and S. C. Liew, "An index policy for minimizing the uncertainty-of-information of Markov sources," *arXiv preprint arXiv:2212.02752*, 2022.
- [35] F. Li, Y. Sang, Z. Liu, B. Li, H. Wu, and B. Ji, "Waiting but not aging: Optimizing information freshness under the pull model," *IEEE/ACM Trans. Netw.*, vol. 29, no. 1, pp. 465–478, 2020.
- [36] C. Papadimitriou and J. Tsitsiklis, "The complexity of optimal queueing network control," in *IEEE CCC*, 1994, pp. 318–322.
- [37] R. R. Weber and G. Weiss, "On an index policy for restless bandits," *Journal of applied probability*, vol. 27, no. 3, pp. 637–648, 1990.
- [38] P. Whittle, "Restless bandits: activity allocation in a changing world," *Journal of Applied Probability*, vol. 25A, pp. 287–298, 1988.
- [39] A. Nedic and A. Ozdaglar, "Subgradient methods in network resource allocation: Rate analysis," in *IEEE CISS*, 2008, pp. 1189–1194.
- [40] I. M. Verloop, "Asymptotically optimal priority policies for indexable and nonindexable restless bandits," 2016.
- [41] B. K. Muirhead, "Mars rovers, past and future," in *IEEE aerospace conference*, vol. 1, 2004.