

AoI-based Finite Horizon Scheduling for Heterogeneous Networked Control Systems

Onur Ayan*, H. Murat Gürsu* Sandra Hirche†, Wolfgang Kellerer*

*Chair of Communication Networks, Technical University of Munich

†Chair of Information-oriented Control, Technical University of Munich

{onur.ayan, murat.guersu, hirche, wolfgang.kellerer}@tum.de

Abstract—Age of information (AoI) measures information freshness at the receiver. AoI may provide insights into quality of service in communication systems. For this reason, it has been used as a cross-layer metric for wireless communication protocols. In this work, we employ AoI to calculate penalty functions for a centralized resource scheduling problem. We consider a single wireless link shared by multiple, heterogeneous control systems where each sub-system has a time-varying packet loss probability. Sub-systems are competing for network resources to improve the accuracy of their remote estimation process. In order to cope with the dynamically changing conditions of the wireless link, we define a finite horizon age-penalty minimization problem and propose a scheduler that takes optimal decisions by looking H slots into the future. The proposed algorithm has a worst-case complexity that grows exponentially with H . However, by narrowing down our search space within the constrained set of actions, we are able to decrease the complexity significantly without losing optimality. On the contrary, we show by simulations that the benefit of increasing H w.r.t. remote state estimation performance diminishes after a certain H value.

I. INTRODUCTION

The next generation communication networks are envisioned to support a wide range of machine-to-machine (M2M) and internet-of-things (IoT) applications, e.g., vehicular networks, industrial automation, smart grid. From system theoretical point of view, they can be modeled as *networked control systems* (NCS) which are feedback control loops that are closed over a communication network.

Studies show that the performance of NCS, i.e., quality of control (QoC), is tightly coupled with the service provided by the communication network. However, the metrics that have been extensively used to measure the performance of human-oriented networks, such as delay, packet loss and throughput, are not sufficient to capture the QoC requirements of heterogeneous NCS applications. Thus, new cross-layer metrics have been adopted for communication protocol design.

Age of information (AoI) is such a relatively new metric that measures information freshness at the receiver monitoring a remote process [1] and used as a cross-layer metric in wireless medium access protocol design [2]–[4]. AoI is defined as the elapsed time since the generation of the freshest information available at the receiver. AoI increases linearly in time, and drops upon a reception of a new update. In such a setting, regardless of its cause, a low update rate at the receiver leads to staleness of the information. Therefore, AoI is a metric combining multiple metrics of communication networks such

as delay and packet loss into a single one. This renders AoI easy to be used as an intermediate metric to calculate control system metrics. In [2], authors consider a multi-user scenario in which each user is prone to a time-invariant packet loss. They formulate a discrete-time wireless scheduling problem and show analytically that the minimum AoI is achieved by updating the user with the highest AoI. They extend their findings in [3] for the equivalent scenario and provide a complexity analysis of the proposed optimal scheduling algorithm together with other sub-optimal, low-complexity algorithms.

AoI has been considered with linear and homogeneous dynamics in time and it cannot be mapped into closed loop control applications directly, where in general the state evolution is a nonlinear function of time and depends on the system parameters. Therefore, the introduction of non-linear aging concept in [5] facilitated the adoption of AoI in NCS research. In the context of NCS, AoI has first been used in [6] where authors employ AoI as an intermediate metric for a resource scheduling problem. They consider a lossless network with limited resources. They show that the proposed heuristic scheduler achieves a higher QoC than the optimal AoI scheduler. The work in [7] follows the same AoI-based derivations from [6] and proposes a cost optimal, threshold-based transmission policy for a single-loop scenario.

The closest work to ours is [8], which studies the problem of scheduling multiple, heterogeneous control systems sharing a common link with packet loss. Similar to the existing literature, they consider the packet loss to be constant over time. They employ the AoI-based metric, i.e., age-penalty, derived in [6] and propose an optimal scheduler that minimizes the discounted age-penalty over an infinite time horizon. However, this solution is not applicable to scenarios with changing channel conditions over time.

To the best of our knowledge both in AoI and NCS research, there is no previous work that considers time-varying channel conditions and at the same time provides optimality in control-aware scheduling. It is a challenging task to guarantee optimality when the network conditions are changing over time, as the scheduling decisions remain optimal only for a certain amount of time.

In this work, we aim to address the open research question of scheduling decisions to guarantee optimal NCS behavior for changing channel conditions. To this end, we study control dependent age-penalty minimization of feedback control

loops that share a wireless communication link with time-varying packet loss probabilities. Time varying conditions in the network prevent the calculation of infinite horizon optimal scheduling policies. On contrary, by applying stochastic optimization and dynamic programming, we propose an online, centralized scheduling policy that is age-penalty optimal for a finite horizon H . The worst-case complexity of our approach grows exponentially but the performance gain expected in return diminishes after a certain H value, hence a finite H value can be selected without incurring too much complexity. To the best of our knowledge, this is the first work proposing an optimal scheduling algorithm for NCS with varying channel conditions.

II. SYSTEM MODEL

We consider a communication network consisting of N independent, linear time-invariant (LTI) control sub-systems. Each sub-system i consists of a plant \mathcal{P}_i , a sensor \mathcal{S}_i and a controller \mathcal{C}_i that are connected through communication. The state of \mathcal{P}_i is observed via an ideal link by \mathcal{S}_i and transmitted in a packet over a wireless link to \mathcal{C}_i . The wireless link is shared among all N sub-systems. We assume each controller-plant pair to be co-located and hence connected through an ideal link. Moreover, each packet contains a single observation.

Time is divided into slots which is also the smallest time unit in our model. We use $t \in \mathbb{N}$ to index time slots. Each packet transmission starting in slot t ends within the same slot. Moreover, medium access is controlled by a centralized scheduler. That is, each sensor \mathcal{S}_i transmits only when it has been allocated a slot by the scheduler. The remaining sensors that are granted access, do not transmit. Suppose an indicator variable $\delta_i(t) \in \{0, 1\}$ that takes the value of 1 if the i -th sub-system is scheduled for transmission in slot t . Then, for any other sensor $j \neq i$, $\delta_j(t) = 0$ holds; or equivalently $\sum_{i=1}^N \delta_i(t) \leq 1$. Any successfully transmitted information in slot t is available at \mathcal{C}_i at $t + 1$.

We consider a time-varying link quality between each sensor-controller pair that affects the loss probability of transmissions. Each packet transmitted at t is lost with probability $p_i(t)$ or received successfully with probability $1 - p_i(t)$. Given that \mathcal{S}_i is scheduled, let $\gamma_i(t) = 1$ indicate a successful transmission with probability $\Pr[\gamma_i(t) = 1 \mid \delta_i(t) = 1] = 1 - p_i(t)$ and $\gamma_i(t) = 0$ a failed transmission with probability $\Pr[\gamma_i(t) = 0 \mid \delta_i(t) = 1] = p_i(t)$, respectively.

We assume block fading where the link quality is approximately constant for a certain number of slots. This period, in which the link quality is assumed to be constant, is called *coherence time*. The scheduler can measure the instantaneous link quality for each sub-system, hence is aware of the loss probabilities $p_i(t)$. However, the coherence time is unknown to the scheduler.

The generation of packets at \mathcal{S}_i is periodic with $D_i \in \mathbb{Z}^+$ slots between two consecutive observations. We call the generation of a packet a *sampling event* and the time between two consecutive sampling events a *sampling period*. Let $t_{i,o}$,




Fig. 1: Evolution of generation time $t_{i,g}(t)$, received time $t_{i,r}(t)$ and the update time $t_{i,u}(t)$ depicted in y-axis versus time in x-axis. $t_{i,g}(t)$ and $t_{i,u}(t)$ are updated periodically with D_i slots, while $t_{i,r}(t)$ can be updated asynchronously. On x-axis with $\delta_i(t_{1,2}) = 1$ two cases of successful packet transmission for sub-system i are depicted. Note that, while $t_{i,r}(t)$ changes subsequent to the reception of any new information, $t_{i,g}(t)$ follows $t_{i,r}(t)$ only at the instances of sampling events.

selected uniformly in $[0, D_i)$, denote the time of the first sampling event of sub-system i ¹. Thus, the set of slots in which a sampling event occurs is given as:

$$\mathcal{G}_i \triangleq \{t_{i,o}, t_{i,o} + D_i, t_{i,o} + 2D_i, \dots\}. \quad (1)$$

Under the assumption that status is Markovian, having received an update, the controller does not benefit from receiving an older observation. Thus, older packets are considered to be obsolete and “non-informative”. Therefore, \mathcal{S}_i discards any older packet upon the arrival of a new update. We introduce a variable $t_{i,g}(t) \in \mathcal{G}_i$ that denotes the generation time of the packet waiting for transmission. Note that, $t \in \mathcal{G}_i$ are the only slots, in which the information at \mathcal{S}_i is updated. Hence, time dynamics of $t_{i,g}$ follows as:

$$t_{i,g}(t+1) = \begin{cases} t+1 & , \text{ if } t+1 \in \mathcal{G}_i \\ t_{i,g}(t) & , \text{ otherwise} \end{cases}. \quad (2)$$

similar to zero-order hold (ZOH) behavior to represent discrete time signals in continuous time. An exemplary evolution of $t_{i,g}(t)$ is depicted in Fig. 1 where a discrete staircase behavior is observed. Additionally, the generation time of the freshest packet received by \mathcal{C}_i until time t is denoted by $t_{i,r}(t)$ with:

$$t_{i,r}(t+1) = \begin{cases} t_{i,g}(t) & , \text{ if } \delta_i(t) \cdot \gamma_i(t) = 1 \\ t_{i,r}(t) & , \text{ otherwise} \end{cases}. \quad (3)$$

In Fig. 1 two successful updates at time instances t_1 and t_2 demonstrate the relationship between $t_{i,r}(t+1)$ and $t_{i,g}(t)$.

¹One can interpret $t_{i,o}$ as the time offset which initiates the operation of sub-systems. As we allow $t_{i,o}$ to be different for each sub-system, we do not enforce any synchronization of sampling among users.

We consider the behavior of the i -th control sub-system is represented by the following LTI model in discrete time:

$$\mathbf{x}_i[k_i + 1] = \mathbf{A}_i \mathbf{x}_i[k_i] + \mathbf{B}_i \mathbf{u}_i[k_i] + \mathbf{w}_i[k_i] \quad (4)$$

with time-invariant system matrix $\mathbf{A}_i \in \mathbb{R}^{n_i \times n_i}$ and input matrix $\mathbf{B}_i \in \mathbb{R}^{n_i \times m_i}$. In addition, $\mathbf{w}_i \in \mathbb{R}^{n_i}$ represents the system noise characterized by a multi-variate Gaussian distribution with zero mean and covariance matrix $\mathbf{\Sigma}_i \in \mathbb{R}^{n_i \times n_i}$, i.e., $\mathbf{w}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_i)$. Time step, k_i , is a counter variable that is initialized with zero and increased by one after each sampling event. Therefore, it indicates in which sampling period a sub-system currently is. Since we assume the sub-systems to operate slower than the network, in slot granularity the evolution of k_i is not strictly increasing. One can find the current k_i of a sub-system at any slot t simply by²:

$$k_i(t) = \left\lceil \frac{t - t_{i,o}}{D_i} \right\rceil. \quad (5)$$

Furthermore, $\mathbf{x}[k_i] \in \mathbb{R}^{n_i}$ and $\mathbf{u}[k_i] \in \mathbb{R}^{m_i}$ denote the system state and control input in sampling period k_i , respectively. At the beginning of each sampling period k_i , i.e., at $t = t_{i,o} + k_i D_i$, the controller \mathcal{C}_i calculates $\mathbf{u}_i[k_i] \in \mathbb{R}^{m_i}$ based on the available observation history. The obtained control input is applied to \mathcal{P}_i during that sampling period. Therefore, we assume that, any state update arriving after the calculation of $\mathbf{u}[k_i]$ is first utilized for the next input $\mathbf{u}_i[k_i + 1]$.

The discrete-time model enables that control and communication progress with different time steps. Namely, in control, changes occur at sampling events³. This is a well-established approach to model digital control systems in the literature [9]. Thereby, it can be argued that from the application's perspective the information of interest ages in discrete steps every D_i slots. Hence, we define AoI as the number of sampling periods elapsed since the generation of the freshest information at \mathcal{C}_i . As a result, the AoI in any slot t is normalized by the sampling period and can be determined as:

$$\Delta_i(t) = \left\lceil \frac{t - t_{i,u}(t)}{D_i} \right\rceil. \quad (6)$$

with:

$$t_{i,u}(t+1) = \begin{cases} t_{i,r}(t+1) & , \text{ if } t+1 \in \mathcal{G}_i \\ t_{i,u}(t) & , \text{ otherwise} \end{cases}. \quad (7)$$

where $t_{i,u}$ denotes the generation time of the latest utilized packet by \mathcal{C}_i . The distinction between $t_{i,r}$ and $t_{i,u}$ is necessary since the controller receives packets during each sampling period that affect neither the control input nor the system state until the next sampling event occurs. In Fig. 1 the update delay of $t_{i,u}(t)$ following the $t_{i,r}(t)$ is demonstrated. $t_{i,u}(t)$ waits for the control system period to update its value. One can imagine

²For $t < t_{i,o}$, $k_i(t)$ takes a negative value. We allow this as $t_{i,o}$ defines the initialization of sub-systems and the system behavior before the operation starts is not taken into account for the remaining analysis.

³We assume that the sampling period of each sub-system is selected small enough such that the changes between consecutive sampling events are negligible.

the incoming information being queued until it is made use of at the end of the sampling period.

It is important to point out that, as the packets are received at least one slot delayed, they can be utilized first to obtain the following control input. Thus, the minimum AoI in our system is one, i.e., $\Delta_i(t) \geq 1, \forall i, t$.

A. Remote Estimation and Control Law

In order to overcome the shortcomings in the sensor-to-controller link, namely the packet loss and limited network resources, we consider an *estimation-based controller* that estimates the plant state remotely. Suppose $\mathbf{x}_i[k_i - \Delta_i[k_i]]$ is the freshest information used by \mathcal{C}_i that is $\Delta_i[k_i]$ sampling periods old. As in [6], it can easily be shown that the conditional expectation of the state which minimizes the mean squared estimation error, i.e., $\hat{\mathbf{x}}_i[k_i] = \mathbb{E}[\mathbf{x}_i[k_i] \mid \Delta_i[k_i], \mathbf{x}_i[k_i - \Delta_i[k_i]]]$ can be obtained from:

$$\hat{\mathbf{x}}_i[k_i] = \mathbf{A}_i^{\Delta_i[k_i]} \mathbf{x}_i[k_i - \Delta_i[k_i]] + \sum_{q=1}^{\Delta_i[k_i]} \mathbf{A}_i^{q-1} \mathbf{B}_i \mathbf{u}_i[k_i - q]. \quad (8)$$

Note that the *estimation-based controller* has to keep track of the last $\Delta_i[k_i]$ control inputs. However, this does not imply any additional communication effort as this information is already present at \mathcal{C}_i . We further assume that each \mathcal{C}_i is aware of the time-invariant system parameters \mathbf{A}_i , \mathbf{B}_i and $\mathbf{\Sigma}_i$.

Moreover, from Eq. (4) and (8) we calculate the estimation error $e_i(t)$ as:

$$e_i[k_i] \triangleq \mathbf{x}_i[k_i] - \hat{\mathbf{x}}_i[k_i] = \sum_{q=1}^{\Delta_i[k_i]} \mathbf{A}_i^{q-1} \mathbf{w}_i[k_i - q]. \quad (9)$$

As a result, the expected mean squared error (MSE) at \mathcal{C}_i is expressed by:

$$\mathbb{E} \left[(e_i[k_i])^T e_i[k_i] \right] = \sum_{r=1}^{\Delta_i[k_i]-1} \text{tr} \left(\left(\mathbf{A}_i^T \right)^r (\mathbf{A}_i)^r \mathbf{\Sigma}_i \right) \triangleq g(\Delta_i[k_i]). \quad (10)$$

The derivations can be found in [6]. It is important to emphasize that the expected MSE is only a function of AoI and independent from the system state $\mathbf{x}_i[k_i]$. Here, we define the $g(\Delta[k_i])$ as a age-penalty function that maps AoI to MSE.

In addition to estimating the system state remotely, the controller determines the control input from control law as:

$$\mathbf{u}_i[k_i] = -\mathbf{L}_i^* \hat{\mathbf{x}}_i[k_i]. \quad (11)$$

where $\mathbf{L}_i^* \in \mathbb{R}^{m_i \times n_i}$ denotes the optimal state feedback gain matrix. \mathbf{L}_i^* is obtained from:

$$\mathbf{L}_i^* = \left(\mathbf{R}_i + \mathbf{B}_i^T \mathbf{P}_i \mathbf{B}_i \right)^{-1} \mathbf{B}_i^T \mathbf{P}_i \mathbf{A}_i, \quad (12)$$

which solves the discrete time algebraic Riccati equation:

$$\mathbf{P}_i = \mathbf{Q}_i + \mathbf{A}_i^T \left(\mathbf{P}_i - \mathbf{P}_i \mathbf{B}_i (\mathbf{R}_i + \mathbf{B}_i^T \mathbf{P}_i \mathbf{B}_i)^{-1} \mathbf{B}_i^T \mathbf{P}_i \right) \mathbf{A}_i. \quad (13)$$

\mathbf{Q}_i and \mathbf{R}_i are weighting matrices of appropriate size that penalize the state and control inputs in the infinite horizon, linear-quadratic-Gaussian (LQG) cost function F_i :

$$F_i = \frac{1}{K} \limsup_{K \rightarrow \infty} \sum_{k_i=0}^{K-1} (\mathbf{x}_i[k_i])^T \mathbf{Q}_i \mathbf{x}_i[k_i] + (\mathbf{u}_i[k_i])^T \mathbf{R}_i \mathbf{u}_i[k_i]. \quad (14)$$

F_i is an indicator of control performance. The lower F_i is, the higher is the *quality of control (QoC)*. From Eq. (12) it is evident that the controller gain \mathbf{L}_i^* depends only on constant system parameters and thus not optimized to take communication parameters into account. Communication-aware controllers exist in the literature, however they are beyond the scope of this paper.

B. Problem Statement

In Eq. (10), AoI appears as the only time-varying parameter. Thus, by controlling the AoI of sub-systems through scheduling, one can control the expected estimation performance. Our goal is to propose a scheduler that maximizes the overall estimation accuracy in the network by minimizing the expected MSE. We are interested in class of scheduling policies π that consist of a sequence of scheduling decisions $\mu(t')$ for the next $H \in \mathbb{Z}^+$ slots, i.e., $\pi = \{\mu(t), \dots, \mu(t+H-1)\}$. H is the *finite horizon* parameter that defines how many future slots is being considered, while taking the current decision. Therefore, the selection of parameter H controls the ‘‘farsightedness’’ of the adopted scheduler.

Each scheduling decision $\mu(t)$ defines the sub-system that is granted medium access where $\mu(t) = i$ and $\mu(t) = \emptyset$ implies $\delta_i(t) = 1$. First, let us define the network state $\mathbf{s} \in \mathbb{Z}^{3 \cdot N}$ as:

$$\mathbf{s}(t) \triangleq [\mathbf{t}_g(t) \quad \mathbf{t}_r(t) \quad \mathbf{t}_u(t)]^T, \quad (15)$$

with:

$$\mathbf{t}_g(t) \triangleq [t_{1,g}(t) \quad \dots \quad t_{N,g}(t)]^T, \quad (16)$$

$$\mathbf{t}_r(t) \triangleq [t_{1,r}(t) \quad \dots \quad t_{N,r}(t)]^T, \quad (17)$$

$$\mathbf{t}_u(t) \triangleq [t_{1,u}(t) \quad \dots \quad t_{N,u}(t)]^T. \quad (18)$$

Next, the set of possible actions given the network state $\mathbf{s}(t)$ is defined as:

$$\mathcal{M}(\mathbf{s}(t)) = \{\emptyset\} \cup \{i : t_{i,g}(t) > t_{i,r}(t)\}. \quad (19)$$

Eq. (19) implies that we constrain the action set and allow sub-systems only with a new information to transmit. Given a state $\mathbf{s}(t)$ and a scheduling decision $\mu(t) = i$, the transition probability to a possible next state \mathbf{s}' strictly depends on the current condition of the corresponding link and is a function of $p_i(t)$. In particular, if \mathbf{s}' is a possible next state after $\mathbf{s}(t)$, then $\Pr[\mathbf{s}(t+1) = \mathbf{s}' \mid \mu(t) = i, \mathbf{s}(t)] \in \{p_i(t), 1 - p_i(t)\}$.

We call the class of scheduling policies mapping states $\mathbf{s}(t)$ into possible scheduling decisions $\mu(t) \in \mathcal{M}(\mathbf{s}(t))$ *admissible* and denote it by Π . Having the class of admissible policies, we can then confine our search within this class of narrowed down set. This enables us to reduce the search space and thereby complexity without losing optimality.

Suppose $C(\mathbf{s}(t))$ denotes the total expected MSE over all sub-systems given the network state at time t as:

$$C(\mathbf{s}(t)) = \sum_{i=1}^N g(\Delta_i[k_i(t)]). \quad (20)$$

Next, we employ $C(\mathbf{s}(t))$ as the state cost and define the expected finite horizon cost $J_\pi(\mathbf{s}(t))$ for any initial state $\mathbf{s}(t)$ and horizon H as:

$$J_\pi(\mathbf{s}(t)) \triangleq \mathbb{E}_\pi \left[\sum_{t'=t}^{t+H} C(\mathbf{s}(t')) \right]. \quad (21)$$

Note the subscript π in J_π and \mathbb{E}_π which indicates that Eq. (21) gives the expected cost when the scheduling policy $\pi = \{\mu(t), \dots, \mu(t+H-1)\}$ is employed over the horizon H . The optimal policy π^* , in which we are interested, is the one that minimizes $J_\pi(\mathbf{s}(t))$; i.e.:

$$J_{\pi^*}(\mathbf{s}(t)) = \min_{\pi \in \Pi} J_\pi(\mathbf{s}(t)) = J^*(\mathbf{s}(t)). \quad (22)$$

Throughout the remaining sections, we refer to the minimization problem in Eq. (22) as the H -stage problem and drop the subscript π for brevity.

III. FINITE HORIZON SCHEDULER

We propose a *finite horizon (FH) scheduler* that takes optimal scheduling decisions for the next H transmission slots. It is shown in [10, p. 25] that for the H -stage problem starting at state $\mathbf{s}(t)$ and time t , the optimal cost $J^*(\mathbf{s}(t))$ can be obtained by minimizing the right side of the Eq. (22) at each stage as:

$$J_{t'}(\mathbf{s}(t')) = \min_{\mu(t') \in \mathcal{M}(\mathbf{s}(t'))} \mathbb{E}[C(\mathbf{s}(t')) + J_{t'+1}(\mathbf{s}(t'+1))], \quad (23)$$

with the terminal cost:

$$J_{t+H}(\mathbf{s}(t+H)) = C(\mathbf{s}(t+H)). \quad (24)$$

The expectation is taken with respect to the total MSE in the network as defined in Eq. (10) and Eq. (20). Otherwise speaking, one can find the optimal cost $J^*(\mathbf{s}(t))$ by iterating backwards in time from stage $H-1$ to stage 0. Consequently, if at each iteration stage, t' , the optimal $\mu^*(t')$ action is taken, which minimizes Eq. (23), the policy $\pi = \{\mu^*(t), \dots, \mu^*(t+H-1)\}$ is optimal achieving the minimum cost $J_{\pi^*}(\mathbf{s}(t)) = J_t(\mathbf{s}(t))$. The proof can be found in [10].

The H -stage problem, as described above, can be modeled as a tree structure, where each node inside the tree represents a network state $\mathbf{s}(t')$ occurring at time t' with $t \leq t' \leq t+H$. The root of the tree corresponds to the current network state $\mathbf{s}(t)$. Moreover, nodes occurring in the same slot form a level together, where the root node constitutes the 0-th level of the tree. Nodes in the last, i.e., H -th, are called *leaf nodes*.

The backwards iteration to solve the H -stage problem can be viewed as visiting all levels of the tree starting from the leaf nodes and taking the optimal action at each level that minimizes the expected cost in the next level. Each node is




Fig. 2: Example 1 level tree structure with 2 sub-systems, i.e., $H = 1$ and $N = 2$. Each edge is labeled with the corresponding transition probability, i.e., $\Pr[\gamma_{\mu(t)}(t) | \mu(t)]$ as the conditional probability the outcome of the scheduled sub-system given the scheduling decision. The state s_1 and s_2 represent the cases where scheduled sub-system successfully updated the information in the controller while the s_3 represent the failure case where none of the sub-system can update the information.

assigned with a cost obtained from Eq. (23). As a result, the operation of the FH scheduler consists of the following steps:

- 1) Initialize the current state $s(t)$ as the root of a tree structure.
- 2) Starting from the root, determine the possible actions at each node, i.e., $\mathcal{M}(t')$ for $\{t' : t \leq t' < t + H\}$ and subsequently all possible next states $s(t' + 1)$ when action $\mu(t')$ is taken.
- 3) Add all possible next states as child nodes to the next level of the tree with the corresponding transition probabilities from the parent node.
- 4) Repeat steps (2)-(4) until the H -th level of the tree is constructed.
- 5) Assign costs to all states starting from the leaf nodes as in Eq. (23) and (24).

For better understanding, consider a toy example with $N = 2$. Suppose sub-systems are sampled every slot, i.e., $D_1 = D_2 = 1$, which is a special case and implies $t_{i,g}(t) = t, \forall i, t$. Fig. 2 illustrates the first level of the tree structure constructed by the FH scheduler. Note that the tree has the root $s(t)$. Outgoing edges are labeled with corresponding transition probabilities. They are written as conditional probabilities, that is, probability of the transmission outcome given the scheduling decision, i.e., $\Pr[\gamma_{\mu(t)}(t) | \mu(t)]$. Thus, the transition probabilities for sub-system i can be written as: $\Pr[\gamma_i(t) = 1 | \mu(t) = i] = 1 - p_i(t)$ and $\Pr[\gamma_i(t) = 0 | \mu(t) = i] = p_i(t)$ where the case of no scheduled subsystem is a deterministic state of $\Pr[\gamma_{1,2}(t) = 0 | \mu(t) = \emptyset] = 1$. Consequently, given the initial state, $\mathbf{s}(t) = [t \ t \ a \ b \ c \ d]^T$, one can construct the next states at time $t + 1$ as follows:

$$\begin{aligned} s_1(t+1) &= [t+1 \ t+1 \ t \ b \ t \ b]^T, \\ s_2(t+1) &= [t+1 \ t+1 \ a \ t \ a \ t]^T, \\ s_3(t+1) &= [t+1 \ t+1 \ a \ b \ a \ b]^T. \end{aligned}$$

For the sub-system with a successful transmission, the $t_{i,r}(t+1)$ and $t_{i,u}(t+1)$ is set to $t_{i,g}(t) = t$ as given in the initial

state $\mathbf{s}(t)$. For those sub-systems with a failed transmission or that are not scheduled, the $t_{i,r}(t+1)$ is kept same as $t_{i,r}(t)$ and $t_{i,u}(t+1)$ is updated to $t_{i,r}(t)$.

Once the complete tree is constructed in a similar fashion, the FH scheduler assigns the cost and best action to each node as defined in Eq. (23). As a result, the optimal decision, $\mu^*(t)$, at current state $s(t)$ is obtained and executed node and executes the optimal decision $\mu^*(t)$ at level 0. It is important to emphasize that, even though the admissible and optimal policy π^* consists of H optimal actions, the steps (1)-(6) have to be repeated every time slot, as the root of the tree changes and only the action at level 0 is extracted from the policy.

The complexity of the FH scheduling algorithm is dominated by the construction of the tree structure. If all sub-systems in the network have a sampling period of 1 slot, which is the worst-case scenario in terms of complexity, the number of states in the tree equals to $(N + 1)^{H+1} - 1$. Depending on the implementation, namely how often each node is visited to obtain the optimal action, the complexity becomes $c \cdot ((N + 1)^{H+1} - 1)$, where c is a positive constant. Nevertheless, in the simplified Big-O notation, c can be omitted and the complexity can be expressed as $\mathcal{O}(N^H)$. Increasing the sampling period further reduces the number of nodes. In particular, not having any update packet to transmit after a successful transmission, as packets are generated only at sampling events, reduces the number of possible actions at some states. In other words, if $t_{i,g}(t) = t_{i,r}(t)$, then $i \notin \mathcal{M}(t)$. Thereby, the amount of nodes in the next level is reduced.

IV. EVALUATION

In order to evaluate the performance of our proposed FH scheduler, we set up a simulation environment with $N = 3$ sub-systems as a minimal scenario to demonstrate the benefit of finite horizon minimization. One can find the most important parts of our simulation code publicly available at: https://github.com/oayan/finite_horizon_scheduling.

We consider each sub-system to have a different system matrices, i.e., $A_{1,2,3} = \{1.0, 1.25, 1.5\}$, representing different task criticalities. The input matrix is given by $B_i = 1.0, \forall i$. Note that, $A_3 = 1.5$ represents the most challenging application. The control sub-systems start with $x_i[0] = w_i[0]$ for all i and system noise is given by $w_i[k] \sim \mathcal{N}(0, 1)$, i.e., $\Sigma_i = 1$. Sub-systems are sampled periodically every 3 slots, i.e., $D_i = 3$. The arrival time of the first sampling event is selected randomly from a discrete uniform distribution as $t_{i,o} = \mathcal{U}\{0, D_i - 1\}$ such that the sub-systems operate in a non-synchronized fashion.

We assume $\mathbf{Q}_i = 1$ and $\mathbf{R}_i = 0, \forall i$. That is, we take only the state cost into account but neglect the penalty for control effort. Therefore, from Eq.(12) the optimal feedback gain matrix is determined as $\mathbf{L}_i^* = A_i$, which corresponds to deadbeat control strategy.

In order to model the time varying nature of the channel, we assume the loss probabilities, i.e., $p_i(t)$, to be normally

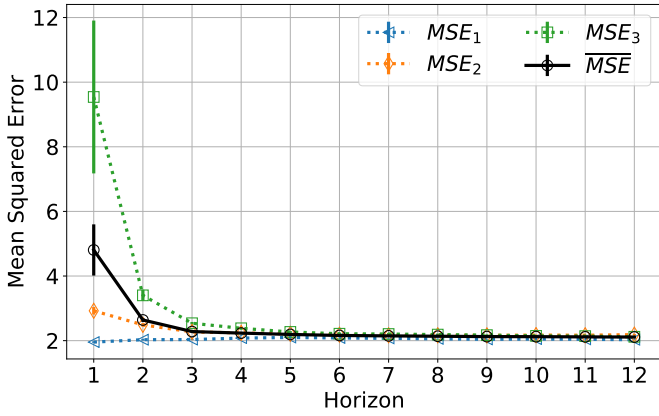


Fig. 3: Solid line \overline{MSE} illustrates the average MSE in the network per time slot. Dashed lines MSE_1 , MSE_2 and MSE_3 show the the average MSE per time slot of each class with $A_1 = 1.0$, $A_2 = 1.25$ and $A_3 = 1.5$, respectively. Vertical bars represent 95% confidence interval.

distributed with mean 0.3 and standard deviation 0.2⁴. The coherence time of the channel, i.e., period of a constant $p_i(t)$, is 30 slots. In other words, the loss probabilities do not vary for 30 slots. The finite horizon scheduler is agnostic to the coherence time and uses $p_i(t)$ available at that time instance.

We vary the finite horizon $H = \{1, \dots, 12\}$ to investigate the impact of making the FH scheduler more far- or shortsighted on the estimation performance. Each configuration is simulated for a duration of $D = 20000$ time slots and repeated $R = 200$ times. We measure the average mean squared error of a sub-system i at each slot, i.e., $MSE_i = \frac{1}{D} \sum_{t=1}^D (e_i(t))^T e_i(t)$. Additionally, the average MSE in the network, i.e., \overline{MSE} , is determined as $\overline{MSE} = \frac{1}{N} \sum_{i=1}^N MSE_i$.

Fig. 3 shows the estimation performance achieved by FH scheduler for varying H . We observe a drastic reduction, i.e., $\sim 45\%$, of \overline{MSE} as we increase the horizon from $H = 1$ to $H = 2$. However, beyond $H = 2$, as the scheduler gets more farsighted, the performance gain diminishes. In particular, the relative decrease in MSE from $H = 5$ to $H = 6$ and from $H = 11$ to $H = 12$ are 1.46% and 0.19%, respectively. It is important to mention that for higher H values, the averages of individual sub-systems, in other words the MSE_i lines, tend to meet. This effect follows from equal weighting of MSE among sub-systems while obtaining the cost in Eq. (20) which in turn leads to equal long-term averages as the scheduling actions become more “long-term optimal”. In addition to MSE, we measure the AoI performance in the network similarly, i.e., $\Delta_i = \frac{1}{D} \sum_{t=1}^D \Delta_i(t)$ and $\overline{\Delta} = \frac{1}{N} \sum_{i=1}^N \Delta_i$.

Fig. 4 illustrates the resulting AoI performance. As we can

⁴The distribution is actually a rectified Gaussian distribution to reflect the possible loss values, i.e., all values below zero and above one are set to 0 and 1, respectively. This is aimed to mimic the packet loss performance in reference 3GPP channel model, where below and above a certain BER just a loss or success is observed, respectively.

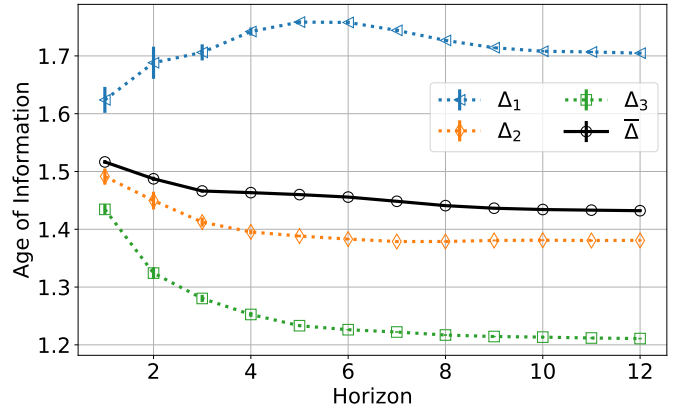


Fig. 4: Solid line $\overline{\Delta}$ illustrates the average AoI in the network per time slot. Dashed lines Δ_1 , Δ_2 and Δ_3 show the the average AoI per time slot of each class with $A_1 = 1.0$, $A_2 = 1.25$ and $A_3 = 1.5$, respectively.

see, the average AoI differs for each sub-system. This is an expected result, since the FH aims to minimize the short-term MSE trajectory, regardless of how often any loop is granted medium access. This leads to unequal distribution of resources due to their heterogeneous system dynamics, thus to different update frequencies. In particular, we can see that the most critical sub-system with $A_i = 1.5$ is scheduled in average more frequently, because in our scenario a lower AoI indicates a more frequent update rate. Additionally, we observe a reduction in Δ_2 and Δ_3 as H increases. This effect results from the increasing foresight of the scheduler for possible high future costs in case of consecutive unsuccessful transmissions. Note that, although Δ_3 is the lowest among all Δ_i , sub-system 3 is still the one with the highest MSE in the network.

From Fig. 4 it is evident that the FH scheduler modifies the distribution of network resources among sub-systems as a different H is selected. On the other hand, Fig. 3 shows that the additional performance gain it brings in terms of cost is not as obvious as the induced complexity, which grows exponentially with H . The effect of increasing complexity is given in table below which compares the average number of nodes in the root tree during simulations (FH) to the worst-case bound without any search space reduction:

H	1	5	9	12
FH	4.6	$5.9 \cdot 10^2$	$5.7 \cdot 10^4$	$1.8 \cdot 10^6$
WC	5.0	$1.4 \cdot 10^3$	$3.5 \cdot 10^5$	$2.2 \cdot 10^7$

We observe a complexity reduction up to approximately 91% in our implementation compared to the worst-case scenario with $\frac{(N+1)^{H+1}-1}{N}$ nodes. This effect is caused by constraining the action set as in Eq. (19) and narrowing down our search space within the class of admissible scheduling policies.

V. CONCLUSION

In this work, together with the system dependent parameters, we propose a novel AoI based control-aware scheduling

algorithm that considers a finite horizon H for a single-hop network with time varying channel conditions. To the best of our knowledge this is the first algorithm that guarantees optimal scheduling decisions with time varying channel conditions. We show by simulations that the quality of control (QoC) can be improved by increasing H since the scheduler becomes more farsighted. On the other hand, the complexity of obtaining such an optimal policy increases exponentially with H . Nevertheless, we show that in spite of the increase in complexity, the QoC does not improve in the same order. In fact, no significant improvement of QoC is observed after $H = 5$ in our setup. Obviously, such a point is strongly scenario dependent and we conclude that one has to find a trade-off between optimality by increasing H and complexity by decreasing H . Most importantly, such a trade-off can be found easily through required modifications to take the dynamics of LTI control systems at hand into account. The proposed scheduler can be deployed in M2M and industrial IoT networks supporting control systems, where the resource management is done centrally.

ACKNOWLEDGEMENT

This work has been carried out with the support of DFG priority programme “Cyber-Physical Networking (CPN)” (SPP 1914).

REFERENCES

- [1] S. Kaul, R. Yates, and M. Gruteser, “Real-time status: How often should one update?” in *2012 Proceedings IEEE INFOCOM*, 2012.
- [2] I. Kadota, E. Uysal-Biyikoglu, R. Singh, and E. Modiano, “Minimizing the age of information in broadcast wireless networks,” in *Allerton Conference on Communication, Control, and Computing*, 2016.
- [3] 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 Transactions on Networking*, 2018.
- [4] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, “Age of information performance of multiaccess strategies with packet management,” *Journal of Communications and Networks*, 2019.
- [5] A. Kosta, N. Pappas, A. Ephremides, and V. Angelakis, “Age and value of information: Non-linear age case,” in *IEEE International Symposium on Information Theory*, 2017.
- [6] O. Ayan, M. Vilgelm, M. Klügel, S. Hirche, and W. Kellerer, “Age-of-information vs. value-of-information scheduling for cellular networked control systems,” in *ACM/IEEE International Conference on Cyber-Physical Systems*, 2019.
- [7] M. Klügel, M. H. Mamduhi, S. Hirche, and W. Kellerer, “Aoi-penalty minimization for networked control systems with packet loss,” in *IEEE INFOCOM WKSHPs*, 2019.
- [8] O. Ayan, M. Vilgelm, and W. Kellerer, “Optimal scheduling for discounted age penalty minimization in multi-loop networked control,” in *IEEE Consumer Communications & Networking Conference*, 2020.
- [9] K. J. Åström and R. Murray, *Feedback Systems: An Introduction for Scientists and Engineers*, 2008.
- [10] D. P. Bertsekas, *Dynamic programming and optimal control*, 4th ed. Athena Scientific, 1995, vol. 1.