

Performance Trade-Offs in Cyber–Physical Control Applications With Multi-Connectivity

Donevski, Igor; Leyva-Mayorga, Israel; Nielsen, Jimmy Jessen; Popovski, Petar

Published in:
Frontiers in Communications and Networks

DOI (link to publication from Publisher):
[10.3389/frcmn.2021.712973](https://doi.org/10.3389/frcmn.2021.712973)

Creative Commons License
CC BY 4.0

Publication date:
2021

Document Version
Publisher's PDF, also known as Version of record

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Donevski, I., Leyva-Mayorga, I., Nielsen, J. J., & Popovski, P. (2021). Performance Trade-Offs in Cyber–Physical Control Applications With Multi-Connectivity. *Frontiers in Communications and Networks*, 2, 1-14. Article 712973. <https://doi.org/10.3389/frcmn.2021.712973>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.



Performance Trade-Offs in Cyber-Physical Control Applications With Multi-Connectivity

Igor Donevski*, Israel Leyva-Mayorga, Jimmy Jessen Nielsen and Petar Popovski

Department of Electronic Systems, Aalborg University, Aalborg, Denmark

OPEN ACCESS

Edited by:

Yuli Yang,
University of Lincoln, United Kingdom

Reviewed by:

Shuai Wang,
Southern University of Science and
Technology, China
Mingzhe Chen,
Princeton University, United States

*Correspondence:

Igor Donevski
igordonevski@es.aau.dk

Specialty section:

This article was submitted to
Communications Theory,
a section of the journal
Frontiers in Communications and
Networks

Received: 21 May 2021

Accepted: 22 July 2021

Published: 16 August 2021

Citation:


Donevski I, Leyva-Mayorga I,
Nielsen JJ and Popovski P (2021)
Performance Trade-Offs in
Cyber-Physical Control Applications
With Multi-Connectivity.
Front. Comms. Net 2:712973.
doi: 10.3389/frcmn.2021.712973

Modern communication devices are often equipped with multiple wireless communication interfaces with diverse characteristics. This enables exploiting a form of multi-connectivity known as interface diversity to provide path diversity with multiple communication interfaces. Interface diversity helps to combat the problems suffered by single-interface systems due to error bursts in the link, which are a consequence of temporal correlation in the wireless channel. The length of an error burst is an essential performance indicator for cyber-physical control applications with periodic traffic, as this defines the period in which the control link is unavailable. However, the available interfaces must be correctly orchestrated to achieve an adequate trade-off between latency, reliability, and energy consumption. This work investigates how the packet error statistics from different interfaces impact the overall latency-reliability characteristics and explores mechanisms to derive adequate interface diversity policies. For this, we model the optimization problem as a partially observable Markov decision process (POMDP), where the state of each interface is determined by a Gilbert-Elliott model whose parameters are estimated based on experimental measurement traces from LTE and Wi-Fi. Our results show that the POMDP approach provides an all-round adaptable solution, whose performance is only 0.1% below the absolute upper bound, dictated by the optimal policy under the impractical assumption of full observability.

Keywords: partially observable Markov decision process (POMDP), interface diversity, multi-connectivity, Gilbert-Elliott, burst error, latency-reliability, Q-MDP

1 INTRODUCTION

In the rise of Industry 4.0, the fourth industrial revolution, there is an amassing interest for reliable wireless remote control operations. Moreover, the application of connected robotics, such as in cyber-physical control, is one of the main drivers for technological innovation toward the sixth generation of mobile networks (Saad et al., 2020). In accord, one of the main use cases for the fifth generation of mobile networks (5G) is ultra-reliable and low-latency communication (URLLC) (Popovski et al., 2019). Reliability and latency requirements for this use case are in the order of $1-10^{-5}$ and of a few milliseconds, respectively. The combination of these two conflicting requirements makes URLLC challenging. For instance, hybrid automatic repeat request (HARQ) retransmission mechanisms provide high reliability but cannot guarantee the stringent latency requirements of URLLC. To solve this, recent 3GPP releases have supported dual- and multi-connectivity, in which data packet duplicates are transmitted simultaneously via two or more paths between a user and a number of eNBs. Hereby, reliability can be improved without sacrificing latency



by utilizing several links pertaining to the same wireless technology—4G or 5G—but at the cost of wasted time–frequency resources (Wolf et al., 2019; Suer et al., 2020b). However, modern wireless communication devices, such as smart phones, usually possess numerous wireless interfaces that can be used to establish an equal number of communication paths. Recent work has proposed interface diversity (Nielsen et al., 2018), which expands the concept of dual- and multi-connectivity to the case where a different technology per interface can be used. Thereby, lower cost connectivity options can help to increase communication reliability. Since constant packet duplication leads to a large waste of resources, the transmission policies in multi-connectivity and interface diversity systems must be carefully designed to meet the performance requirements while avoiding resource wastage and over-provisioning. Furthermore, as we will observe in *Results*, acquiring sufficient knowledge on the channel statistics is essential to attain adequate trade-offs between resource efficiency and reliability in interface diversity systems.

From its definition, the URLLC use case treats each packet individually and, hence, does not capture the performance requirements of numerous applications. For instance, the operation of cyber–physical control applications, which transmit updates of an ongoing process, is usually not affected by individual packets that violate the latency requirements (i.e., untimely packets). Instead, these applications define a *survival time*: the time that the system is able to operate without a required message (3GPP, 2021). Hence, the reliability of communication in such cyber–physical systems is defined by the statistics of consecutive untimely packets, that is, the length of error bursts. Hence, in cyber–physical systems, having multiple interfaces with diverse characteristics is greatly valuable, as it allows to select the appropriate interface based on the requirements of the task at hand. For example, while an LTE-based system with multi-connectivity capabilities and selective packet duplication could satisfy the requirements of the application, it seems likely that a combination of unlicensed (e.g., Wi-Fi) and licensed (e.g., LTE) technologies could lead to similar performance guarantees while achieving a lower usage of scarce licensed spectrum and reduce overall costs.

In this paper, we therefore study the performance of interface diversity in terms of burst error distribution in a

source–destination system, where we consider two fundamentally different technologies: 1) LTE, which is based on orthogonal frequency division multiple access (OFDMA), which operates in licensed spectrum, and where the base station (BS) schedules the uplink resources for communication, and 2) Wi-Fi, which is based on carrier sense multiple access (CSMA) and operates in unlicensed spectrum. The goal of the proposed interface diversity system, as illustrated in **Figure 1**, is addressing the survival time in cyber–physical control applications (3GPP, 2021). In particular, we investigate the trade-offs between system lifetime (the time until the system reaches the end of the survival time and operation is interrupted) and energy consumption. Given the nature of periodic traffic, the survival time can be expressed as the maximum tolerable number of consecutively lost or untimely information packets. To consider the effect of channel correlation in consecutive errors, we use the Gilbert–Elliott (Haslinger and Hohlfeld, 2008) model that is well suited for representing time-correlated transmissions (Willig et al., 2002). Using this approach, we formulate the problem as a partially observable Markov decision process (POMDP) that takes into account the limited observability of the inactive interfaces. Hence, based on the observations and the belief states, we can calculate the optimal transmission policy even for devices with extremely limited computational power. We observed that the performance trade-offs achieved with the POMDP approach are greatly similar when compared with the ones achieved with an idealized fully observable MDP. The key contributions of this work are as follows:

- An interface diversity problem is formulated for energy-constrained devices as a POMDP. Hence, our approach considers the limited observability of the inactive interfaces: those that do not transmit and thus do not receive feedback. While our results are presented for a device using an LTE and a Wi-Fi interface, our model is sufficiently general and, hence, can be applied to cases with more than two interfaces and to different technologies.
- Interface diversity policies are analyzed for cyber–physical control applications, where a certain number of untimely packets are tolerated and those with error burst due to the temporal correlation in the wireless channel are considered by means of a Gilbert–Elliott model.

- A computationally simple solution, the Q-MDP value method, can be used for solving the POMDP. Using this method, we obtain results that closely follow the performance of the fully observable MDP. Specifically, the expected loss in the reward is only around 0.1%.

The rest of this paper is organized as follows. We initially present an elaborate explanation of multi-connectivity and interface diversity's role in timeliness in **Section 2**. Next, we present the system model in **Section 3**, followed by the analysis of the scenario and our proposed method to solve the POMDP in **Section 4**. Then, we present the numerical results in **Section 5**. Finally, we conclude this paper with a summary of the work in **Section 6**.

2 LITERATURE REVIEW

Multi-connectivity has been studied from different perspectives. For instance, Wolf et al. (2019) studied a scenario with one user equipment (UE) connected to multiple BSs and with multiple simultaneous connections to the same BS. The benefits of this approach are assessed in terms of transmit power reduction, achieved by increasing the signal-to-noise ratio (SNR). Following a similar multi-connectivity approach, a matching problem is formulated by Simsek et al. (2019), where the number of UE in the network and the limited wireless resources are considered. The objective is to provide the desired reliability to numerous users by assigning only the necessary amount of resources to each of them. Mahmood et al. (2018) investigated a similar problem in a heterogeneous network scenario with a small cell and a macrocell. Their results show that multi-connectivity is particularly useful for cell-edge UE connected to the small cell and provides even greater benefits when URLLC and enhanced mobile broadband (eMBB) traffic coexist. She et al. (2018) considered multi-connectivity for URLLC as a combination of device-to-device and cellular links, where correlated shadowing is considered. They achieved remarkable increases in the availability ranges for both interfaces. In our previous work (Nielsen et al., 2018), we studied the benefits of interface diversity in terms of reliability for a given error probability. Finally, Suer et al. (2020a) assessed the performance scheduling schemes such as packet duplication and load balancing in order to achieve latency and reliability improvements. The authors exploited a combination of a local Wi-Fi and a private LTE network, which was tested under traffic patterns that are expected to appear in an industrial communication setting.

In the studies mentioned above, only stationary error probabilities are considered. Moreover, Dzung et al. (2015) provided a thorough investigation of switching off a singular interface that has an unreliable channel, based on channel feedback. The goal of the authors is thus aligned with ours since they aim for an energy-efficient transmission policy given bursty channels, for reliable connectivity of synchronous services. However, the use of different interfaces provides unique benefits for URLLC, especially in the case of bursty wireless errors. For instance, different interfaces are likely to present

different burst error distributions, and the correlation of errors between different interfaces is expected to be much lower compared to the correlation between multiple links using the same wireless interface. Despite these evident benefits, and the thorough investigation of burst errors in past research (Yajnik et al., 1999), little research has been conducted on interface diversity with error bursts. Specifically, our previous work presents one of the few analyses of this kind (Nielsen et al., 2019). However, it was limited to the benefits of interface diversity in the length of error and success bursts without considering the impact on resource efficiency.

Cyber-physical control applications can belong to one of two major categories depending on the traffic direction requirements: in downlink or uplink only (open-loop control) or the combined uplink and downlink (closed-loop control) requirements (3GPP, 2021). Moreover, a closed-loop control application needs to process incoming events and thus give appropriate instruction commands to those events (Ploplys et al., 2004). In such scenarios, timeliness is critical to avoid violating the system's imposed latency requirements, which leads to executing outdated actions. Therefore, being untimely is the equivalent of a failure in communication service availability.

Open-loop control applications with periodic commands appear frequently in industrial applications and are considered representative of cyber-physical control systems. In these applications, failing a specific number of consecutive updates directly corresponds to exceeding the survival time and, hence, to an error in the system. For example, it has been observed that the number of consecutive errors impacts the stability of the system and leads to a considerable decrease in safety of autonomous guided vehicles (de Sant Ana et al., 2020). As in the present model and in our previous work (Nielsen et al., 2019), a Gilbert-Elliott model was considered by de Sant Ana et al. (2020) to introduce correlation in the wireless channel. Finally, the novelty of this work comes from investigating the problem of interface diversity for timely packet arrivals for cyber-physical control applications in a burst error channel, where the reliability of the system comes as a trade-off of energy.

3 SYSTEM MODEL

We consider a point-to-point communication between a user and a BS in an industrial scenario. The user samples a given set of physical phenomena and generates data periodically, where T_s is the sampling period. The sampled data are immediately transmitted to the BS, where they are used for control purposes, so that they must be received within a pre-defined latency constraint $\theta \leq T_s$. Hence, it is now convenient to introduce the definition of the latency-reliability function, which stands for the probability of being able to transmit a data packet from a source to a destination with a given latency deadline (Nielsen et al., 2018).

Let L be the RV that defines the packet latency. Then, for a given interface i and latency deadline θ , the latency-reliability function is defined as

$$F_i(\theta) = \Pr(L \leq \theta \mid i). \quad (1)$$

As such, the latency-reliability function is a CDF of the interface's latency, where lost packets have the equivalent of




FIGURE 2 | (A) Two-state GE model for an interface i and **(B)** four-state GE model for a user with two interfaces.

infinite latency. Thus, the error probability becomes a specific value (of deadline) Θ in the latency–reliability function, and we define the probability of error for an interface i as

$$P_e^{(i)} = 1 - F_i(\theta). \quad (2)$$

It should be noted that the traditional definition of the probability of error is obtained for the case $\theta \rightarrow \infty$ and that the distribution of L can be updated continuously to reflect the changes in the wireless channel.

We consider the case where the interface diversity system uses packet cloning, where a full packet is transmitted via each of the N interfaces. Next, by assuming that errors across the multiple available interfaces occur independently, the end-to-end error probability can be calculated as (Billinton and Allan, 1992; Nielsen et al., 2018)

$$P_e^{\text{E2E}} = \prod_{i=1}^N (1 - F_i(\theta)) = \prod_{i=1}^N P_e^{(i)}. \quad (3)$$

Note that the correlation of the large-scale fading across the interfaces is captured by the model described above through the distribution of the RV L . The assumption of errors occurring independently across interfaces holds since correlation in the fast fading may only occur if the antenna elements within an array have insufficient spacing and/or if the concurrent transmissions occur in frequencies that are separated by less than one coherence bandwidth (Chen et al., 2021). In our case, the use of two different technologies and frequency bands for Wi-Fi (unlicensed ISM bands) and LTE (licensed spectrum) ensures that the transmissions are sufficiently separated in frequency to avoid correlation.

The BS sends individual feedback per interface to the user after each transmission attempt. If the data are not received within θ , they are declared as missing and the user receives a NACK. The system tolerates a maximum number of missed transmissions. Specifically, if the number of missed transmissions is N , the system declares a failure and operation is interrupted. Otherwise, the system is able to continue normal operation whenever the number of missed transmissions is $n \leq N$ (smaller than the survival time).

In the following, we define our interface diversity problem as a POMDP denoted as the tuple $(\mathcal{S}, \mathcal{A}, T, \mathcal{R}, \Omega, \mathcal{O})$. Here, \mathcal{S} is the set of states, \mathcal{A} is the set of actions, T is the transition probability to the next state of the environment given a state–action pair, $\mathcal{R} \subset \mathbb{R}$ is the set of

immediate rewards, Ω is the set of possible observations, and \mathcal{O} is the observation probability when transitioning.

3.1 The Environment

The user, i.e., the agent, interacts with the *environment* at discrete time steps $t \in \mathbb{N}$ by sampling and transmitting data to the BS. It is equipped with two distinct communication interfaces, $i \in \{1, 2\}$, where the generated data can be transmitted. Throughout this paper, we assume that these interfaces are completely independent from each other and that each interface is accurately modeled by a two-state Gilbert–Elliott (GE) model (Haßlinger and Hohlfeld, 2008). The GE model was selected due to its simplicity and to the ability to capture temporal correlation.

In our GE model, the details of the implemented protocol and the wireless conditions—interference, noise, and fading—are simplified and related to two possible states in a discrete-time Markov chain (DTMC). These are the *good* state G and the *bad* state B . Hence, we model the state space for the GE model for an interface i as $\mathcal{S}_i = \{G, B\}$. At any given point in time, an interface is in the G state if the protocol and the wireless channel conditions are such that allowing for a transmission to be received within the latency constraint θ . Otherwise, the interface is in the B state.

This simple GE model has two parameters, namely, p_i and r_i , which determine the transition probabilities and, hence, the steady-state error probability and burst lengths (Haßlinger and Hohlfeld, 2008). Hence, these are system- and environment-specific and can only be learned after deployment by collecting statistics of the packet transmissions. An additional benefit of using the GE model is that, through continuous tuning of the statistical parameters, it allows to capture cross-interface correlation due to large-scale fading or traffic surges.

We denote the state of interface i at time step $t - 1$ as s_i and as s'_i at time step t . The parameter p_i represents a transition from state G to B and r_i from B to G (i.e., a recovery from the bad state). Hence, the transition probabilities are defined as

$$\Pr(s'_i = G | s_i = G) = 1 - p_i, \quad (4)$$

$$\Pr(s'_i = G | s_i = B) = r_i, \quad (5)$$

$$\Pr(s'_i = B | s_i = G) = p_i, \quad (6)$$




FIGURE 3 | One-step transitions from an arbitrary non-absorbing state (i.e., $n < N$) for the partially observable Markov decision process (POMDP) with two interfaces.

$$\Pr(s'_i = B \mid s_i = B) = 1 - r_i. \quad (7)$$

Figure 2A illustrates the GE model with one interface, whose transition probability in a matrix form is

$$\mathbf{P}_i = \begin{bmatrix} 1 - p_i & p_i \\ r_i & 1 - r_i \end{bmatrix}. \quad (8)$$

Building on this, the state of a system with two interfaces is defined by the four-state GE model illustrated in **Figure 2B**, where transition labels are omitted for brevity. To elaborate, transition probabilities are calculated under the assumption of the two interfaces being independent, for example, the transition from state G,G to state B,B has probability $p_1 p_2$.

Besides the status of each interface, knowing the number of consecutive missed data n is essential for the operation of the system. Therefore, to build a respective Markov decision process (MDP), we define the state space as $\mathcal{S} = \{(s_1, s_2, n)\}$, where $n \in \{0, 1, 2, \dots, N\}$. Note that the true state of the interface is observable only after concluding the transmission at time step $t - 1$, and before the following transmission attempt at time t . Thus, the state of the environment at time t denoted as $S \in \mathcal{S}$ is defined by the outcome of the transmission in the last attempt at $t - 1$. Hence, all states $S \in \mathcal{S}$: $n = 0$ indicate that a transmission at time $t - 1$ was received successfully. Furthermore, all states $S \in \mathcal{S}$: $n = N$ are absorbing states, and consequently, this is a finite MDP with episodic tasks.

Ultimately, the goal of the system is to navigate the MDP in a way that decreases the amount of errors and altogether reduces the likelihood of having N consecutive errors. Thus, the system should be incentivized to maximize its expected lifetime while optimizing the costs associated with each transmission. This is done through proper allocation of rewards for each action in the MDP. However, the main challenge for solving the issue comes as a product of the limited observability of the defined MDP system

when an interface is switched off. The details for this are encompassed in the following subsection.

3.2 Actions, Rewards, and Uncertainty

At each time step t (i.e., data transmission instant), the user takes an action $A \in \mathcal{A} = \{(a_1, a_2)\}$, where $a_i \in \{0, 1\}$; $a_i = 1$ indicates transmission and $a_i = 0$ indicates no transmission for an interface i . Hereafter, we denote that the interface i is on when $a_i = 1$ or off otherwise. Note that, in our case, $\mathcal{A}(S) = \mathcal{A}$ for all $S \in \mathcal{S}$; that is, the set of actions is the same in every possible state. This totals to three different actions—leaving out the option to turn off all interfaces altogether $A = (0, 0)$, having either interface off, or having both interfaces turned on.

Having taken action A when in state S , there is a probability $T(S, A, S')$ to end up in state S' ; therefore, it must apply that $\sum_{S'} T(S, A, S') = 1$. Note that S' represents the true state of both interfaces at time t , which is revealed only after taking action A . A missed transmission can thus occur when both interfaces are transmitting $A = (1, 1)$ but are in the bad state $S' = (B, B, n)$, or a single interface is transmitting that is in its corresponding bad state— $A = (0, 1)$ when $S' = (G, B, n)$ or $A = (1, 0)$ when $S' = (B, G, n)$.

Therefore, following a transmission action A , the user receives a reward R that also depends on the feedback by the BS given before time $t + 1$. Having arrived at state S' by taking action A when in state S yields a reward $R(S, A, S')$ that is a function $r(n)$, where $r(n = 0) = 1$ is a successful transmission and $r(n > 0) = -1$ is a missed transmission. In the overall reward allocation, we also account for the cost of using the interface i , specifically,

$$R(S, A, S') = r(n) - c(A) = r(n) - a_1 c_1 - a_2 c_2, \quad (9)$$

where $c(A)$ is the cost of taking action A .

In an MDP, a policy π is a function that maps each state $S \in \mathcal{S}$ to an action $A \in \mathcal{A}(s)$. Therefore, given a policy $\pi(S)$, the agent will choose action A , once it finds itself in state S . Our objective is to find an optimal policy π^* that selects the best action given some




FIGURE 4 | An illustration of a single step of the value iteration process for state S that tests the Q -function for all state–action pairs, where future states follow the optimal policy π^* .

state S . Let K denote the total number of time steps until the system transitions to an absorbing state. The best action at time t is the one that maximizes the discounted return.

$$G_t = \sum_{k=0}^{K-t-1} \gamma^k R_{t+k+1}, \quad (10)$$

where $0 < \gamma < 1$ is the discount factor, and the system lifetime is $K < \infty$ as there is always a non-zero probability of transitioning to the absorbing states in a finite number of steps. Thus, the value of the MDP, when starting from initial state S_{init} under a policy π , is

$$V_{\pi}(s) = \mathbb{E}_{\pi}[G | S = S_{\text{init}}]. \quad (11)$$

Next, we define the set of observations as $\Omega = \{(o_1, o_2, n)\}$, where $o_i \in S_i \cup \{0\}$ is the set of observations for an interface i , and the number of consecutive missed deadlines n is always observable. The transition from S to S' following action A provides a deterministic observation $O = (o_1, o_2, n) \in \Omega$ in the following manner: (Since all observations are deterministic with each action, the set \mathcal{O} is irrelevant to the analysis.) I) Having an interface on, namely, $a_i = 1$, allows for fully observing the state of that interface $o_i = s_i$. II) On the contrary, having an interface off, namely, $a_i = 0$, provides no observation of the state of that interface, i.e., $o_i = 0$, unless additional mechanisms are available to correctly estimate the state, for example, based on the exchange of control messages. Building on this, we define the following function:

$$o_i = \begin{cases} s_i & \text{if } a_i = 1 \text{ or } \alpha = 1, \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

where $\alpha = 1$ indicates that the BS has additional mechanisms to perform the observation. **Figure 3** illustrates the action space along with the associated transitions to state S' and observations from an arbitrary non-absorbing state S .

4 ANALYSIS

As a first step, we derive the steady-state probabilities of the good and bad states from the transition matrix \mathbf{P}_i as (Haßlinger and Hohlfeld, 2008)

$$\pi_{i,G} = \frac{r_i}{p_i + r_i}, \quad (13)$$

$$\pi_{i,B} = \frac{p_i}{r_i + p_i}, \quad (14)$$

where $\pi_{i,G} + \pi_{i,B} = 1$.

Assuming that the system has both interfaces turned on during initialization, the initial state of the MDP S_{init} is chosen randomly among the set of states $\{(G, G, 0), (G, B, 0), (B, G, 0), (B, B, 1)\}$ based on the steady-state probabilities $\pi_{i,G}$ and $\pi_{i,B}$.

4.1 Policy Utility Through the Value and Q-Functions

Let $V_{\pi}(S)$ be the expected utility received by following policy π from state S :

$$V_{\pi}(S) = \begin{cases} Q_{\pi}(S, A) & \text{if } S \notin \{(s_1, s_2, N)\}, \\ 0 & \text{otherwise,} \end{cases} \quad (15)$$

where $Q_{\pi}(S, A)$ is the expected utility of taking action A from state S and then following policy π (Puterman, 2014). Starting from state S_{init} , our goal is to find the optimal policy π^* that results in the maximum value that can be obtained through any policy $V_{\pi^*}(S_{\text{init}})$. As illustrated in **Figure 4**, when not following the optimal policy π^* , but sampling the value of the Q -functions for each action A in the space of values $V_{\pi^*}(S')$, we get

$$Q_{\pi^*}(S, A) = \sum_{S'} T(S, A, S') [R(S, A, S') + \gamma V_{\pi^*}(S')], \quad (16)$$

where γ is the discount factor that controls the importance of short-term rewards (γ values close to 0), or long-term rewards (γ values close to 1), where the anticipated rewards are represented through the value of the state as

$$V_{\pi^*}(S) = \begin{cases} \max_{A \in \mathcal{A}} Q_{\pi^*}(S, A) & \text{if } S \notin \{(s_1, s_2, N)\}, \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

Unless the user is in an N state that is absorbing, the values $V_{\pi^*}(S)$ are recurring and can be approximated through the iterative process of (Puterman, 2014)

$$V_{\pi^*}^k(S) \leftarrow \max_{A \in \mathcal{A}} \sum_{S'} T(S, A, S') [R(S, A, S') + \gamma V_{\pi^*}^{k-1}(S')], \quad (18)$$

which continues until it converges to some predefined precision ϵ of the past and current values:

$$\max_{S \in \mathcal{S}} |V_{\pi^*}^k(S) - V_{\pi^*}^{k-1}(S)| \leq \epsilon. \quad (19)$$

Thus, the method of value iteration guarantees finding the optimal value for an MDP that is a function of the optimal policy. Given a κ number of iterations to converge to a solution, the complexity of this algorithm is $O(\kappa S A S')$, which given our small MDP is insignificant. Unfortunately, this method does not directly produce an optimal policy for a POMDP; however, this can be addressed by the Q_{MDP} value method.

4.2 Belief Averaged Q_{MDP} Value Method

Due to the limited information on the channel properties for each interface, our Markov process is a POMDP where we cannot fully

observe the true state space for time t . Therefore, the agent maintains a belief b on the state of the system $S \in \mathcal{S}$ based on the observation $O \in \Omega$. By observing that there is no uncertainty on the value of n and that the state of the interfaces \mathcal{S}_i is independent, we can define the belief as

$$b(S, O) = \Pr(S = (s_1, s_2, n) \mid O = (o_1, o_2, n)) = b_1(s_1, o_1)b_2(s_2, o_2), \quad (20)$$

where b_i is the belief for an interface i to be in state s_i given its observation o_i . The b_i values are updated recursively with each following observation as

$$b_i(s_i, o_i) \leftarrow \Pr(s_i \mid o_i) = \begin{cases} 1 & \text{if } s_i = o_i \text{ and } o_i \neq 0, \\ f_{s_i}(b_i) & \text{if } o_i = 0, \\ 0 & \text{otherwise,} \end{cases} \quad (21)$$

where $f_{s_i}(b_i)$ is a function for calculating the probability of being in state s_i as a function of the previously held beliefs:

$$f_G(b_i) = (1 - p_i)b_i(G, 0) + r_i b_i(B, 0), \quad (22)$$

$$f_B(b_i) = p_i b_i(G, 0) + (1 - r_i)b_i(B, 0). \quad (23)$$

Therefore, knowing the belief $b(S, O) \forall S \in \mathcal{S}$, we can proceed with finding an optimal policy for the underlying POMDP through the following two steps:

Step 1: Ignore the observation model and compute the Q-values $Q_{\pi^*}(S, A)$ given directly from the state-action pairs. These are denoted as $Q_{MDP}(S, A)$ and are obtained through calculating the Bellman operator in the value iteration method.

Step 2: Calculate the belief averaged Q-values for each action and belief $b(S, O)$ as

$$Q_A(b) = \sum_{S \in \mathcal{S}} b(S, O) Q_{MDP}(S, A). \quad (24)$$

The optimal policy now becomes a function of the belief, instead of the current state, and is

$$\pi^*(b) = \max_{A \in \mathcal{A}} Q_A(b). \quad (25)$$

Note that this is a method that does not incentivize updating the belief state but optimizes with the assumption that we will have full observability following the transmission at time t (Littman et al., 1995).

4.3 Parameter Tuning

For the MDP to optimize the operation of the underlying communication system, we require a proper assignment of the rewards and costs for the MDP. Therefore, the reward and punishment for a successful or a missed transmission were fixed to 1 and -1, respectively. Conversely, the value of c_i , the cost of using the interface i , greatly depends on the specific characteristics of the system and on the individual notion of resource efficiency. Moreover, the cost of using an interface is directly related to the consumption of resources that would otherwise be available to other services. Throughout the rest of this paper, we consider that the cost of using an interface is given by the energy consumption. However, other parameters can be




FIGURE 5 | Empirical latency CDFs of considered interfaces.

used to define the cost of each interface when adapting our methods to a specific system.

Given a transmission power E_{LTE} and E_{Wi-Fi} for the LTE and the Wi-Fi interface, respectively, we calculate the cost for an interface i as

$$c_i = \eta \frac{E_i}{\sum_i E_i}, \quad (26)$$

where η is a cost scaling factor that serves to reduce/increase the importance of the energy transmission costs with regard to the initial rewards. The scaling factor η was sampled across several values in the range $0 \leq \eta \leq 1$, which resulted in five different optimal policies, one for each different $\eta \in \{0, 0.03, 0.07, 0.2, 1\}$.

4.4 Latency Measurements for Modeling Wi-Fi and LTE

Traces of latency measurements for different communication technologies were obtained by sending small (128 bytes) UDP packets every 100 ms between a pair of GPS time-synchronized devices through the considered interface (LTE or Wi-Fi) during the course of a few work days at Aalborg University campus. A statistical perspective of these data is given by the latency CDFs in Figure 5, which clearly outlines some key differences between the performance of the LTE and Wi-Fi interfaces. While Wi-Fi can achieve down to 5 ms one-way uplink latency for 90% of packets, it needs approx. 80 ms to guarantee delivery of 99% of packets. For LTE, on the contrary, there is hardly any difference between the latency of 90 and 99% delivery rates, approx. 36 and 40 ms, respectively. Since the measurements for both LTE and Wi-Fi were recorded in good high-SNR radio conditions, we expect that the differences between LTE and Wi-Fi can, to a large extent, be attributed to the inherent differences in the protocol operation and the fact that LTE operates in the licensed spectrum, whereas Wi-Fi has to contend for spectrum access in the unlicensed spectrum.

4.5 Performance Evaluation

To conduct the performance evaluation of the policies obtained with the POMDP, we define the following benchmarks:

- A *fully observable system* assumes an inherent ability of the BS to inform the user about the interface that is turned off, for example, by using pilots that precede the transmissions. In this case, $\alpha = 1$, making the POMDP collapse to an MDP. We denote the policy with full observability as $\pi_{\alpha=1}^*$.
- The *forgetful POMDP (F-POMDP)* maintains a single state of partial belief and, afterward, assumes the steady-state probabilities $\pi_{i,G}$ and $\pi_{i,B}$ for the inactive interface. This forgetful approach collapses to a small MDP where belief does not need to be continuously computed.
- The *hidden MDP (H-MDP)* is the fully reduced MDP of the forgetful approach, where the belief averages in F-POMDP are joint in a single state. Here, the transition probabilities for the inactive interface directly become the steady-state probabilities $\pi_{i,G}$ and $\pi_{i,B}$.

The obtained policies are evaluated based on the following performance indicators: i) the distribution of the number of consecutive errors n ; ii) the utilization of the LTE interface, defined as the ratio of time slots when the LTE interface is turned on $u_{\text{LTE}}(\pi) = \Pr(a_1 = 1 \mid \pi)$; iii) the expected system lifetime, defined as the number of time steps from initialization until the system transitions into an absorbing state. For the latter, let $K(\pi)$ be the expected system lifetime with policy π ; and iv) the expected total reward of the system with policy π , from initialization until absorption, defined as $\bar{R}(\pi)$.

Building on this, we assess the policies derived with partial observability w.r.t. the policy with full observability based on the following:

- *System lifetime delta* denotes the relative increase of the expected system lifetime w.r.t. the MDP with $\alpha = 1$ (i.e., full observability), defined as

$$\Delta K = \frac{(\bar{K}(\pi) - \bar{K}(\pi_{\alpha=1}^*))}{\bar{K}(\pi_{\alpha=1}^*)}. \quad (27)$$

Hence, positive values of ΔK indicate an increase in the system lifetime w.r.t. the optimal policy with full observability.

- *Policy deviation* measures the relative change in the expected system lifetime K and expected transmission cost as

$$\Delta \pi = |\Delta K| + |u_{\text{LTE}}(\pi) - u_{\text{LTE}}(\pi_{\alpha=1}^*)| \cdot c_{\text{LTE}}. \quad (28)$$

Note that this measures the difference in behavior w.r.t. the optimal policy but does not necessarily reflect a proportional decrease in performance. Instead, this is a measure of the normalized collective error, as in common estimators that try to project the optimal LTE usage and system lifetime.

- *Relative reward loss* defines the relative loss in the expected total reward $\bar{R}(\pi)$ with policy π w.r.t. $\pi_{\alpha=1}^*$ as

TABLE 1 | Parameters for evaluation.

Label	Definition	Value
θ	Latency constraint	38.25 ms
ρ_{LTE}	LTE's p-transition probability	0.0178
$\rho_{\text{Wi-Fi}}$	Wi-Fi's p-transition probability	0.0515
r_{LTE}	LTE's r-transition probability	0.2577
$r_{\text{Wi-Fi}}$	Wi-Fi's r-transition probability	0.9468
$E_{\text{Wi-Fi}}$	Power consumption of the Wi-Fi interface	15.85 mW
E_{LTE}	Power consumption of the LTE interface	200 mW
N	Maximum number of consecutive missed transmissions	4
ϵ	Value iteration convergence criteria	10^{-11}
k_{max}	Value iteration maximum number of iterations	10^5
γ	Discount factor	0.99999

$$\mathcal{L}(\pi) = \frac{|\bar{R}(\pi_{\alpha=1}^*) - \bar{R}(\pi)|}{\bar{R}(\pi_{\alpha=1}^*)}. \quad (29)$$

The results with the fully observable MDP were obtained analytically. In order to evaluate the performance of the POMDP and forgetful methods, analytical results were obtained for extreme values of parameter η . For all other cases, we performed Monte Carlo simulations of 20,000 episodes. The duration of each episode depends on the system lifetime which could last up to several million time steps.


5 RESULTS

In this section, we investigate the performance of the modeled system. The investigation in this section is guided by the use of interface diversity in the case of a combination of Wi-Fi and LTE. The performance of the aforementioned system where $i = 1$ is LTE and $i = 2$ is Wi-Fi was evaluated by a Monte Carlo Matlab simulation (when necessary) where the calculation of the statistical properties for the GE model is derived from experimental latency measurements.

We tested the system for all five different values $\eta = 0, 0.03, 0.07, 0.2, 1$ where the fully observable MDP system had different transmission policies. Given the measurements and the characteristics of our measurement setup, we tuned the simulation parameters to the values in **Table 1**.

5.1 Extreme Policies

As a starting point, we describe and evaluate the policies obtained in the cases where the value of parameter $\eta \in (-\infty, \infty)$ is set to an extremely low or high value (i.e., at either of the extremes of its range). When adjusting the scaling factor to its lowest possible value $\eta \rightarrow -\infty$, the cost of using each interface is omitted and the MDP optimizes in favor of not losing any transmissions. Thus, we observe an extreme behavior that is not affected by the belief, or the POMDP behavior. Specifically, the optimal policy maintains both interfaces turned on no matter the current state $\pi^*(S) = A(1, 1) \forall S \in \mathcal{S}$. Since the utilization of both interfaces is 100%, this gives the upper bound on burst




error performance for the whole system. Even in this case, there is a non-zero probability to end up in the absorbing state which happens with an expectation of $5.0738 \cdot 10^6$ transitions. This is the defined lifetime of the system, and before turning off due to failure, the system maintains successful transmissions for 99.67% of the time, 0.32% of transmissions have a single error, 0.0127% have two consecutive errors, and $4.9971 \cdot 10^{-4}\%$ of all burst errors have three consecutive errors.

On the contrary, setting $\eta \rightarrow \infty$ creates a lower bound on the system performance that aims to minimize the cost of operation at the expense of decreasing the system lifetime. This is the result of scaling the cost of performing a transmission to be higher than the reward of maintaining successful transmissions. Since the action space is restricted to use at least one interface for transmission at all times, in such a cost-restraint system, it is reasonable to only allow for the utilization of the Wi-Fi interface $\pi^*(S) = A(0, 1) \forall S \in \mathcal{S}$, due to the high cost of using LTE. Since we have 100% utilization of Wi-Fi and 0% utilization of LTE, the system lifetime of the system decreases drastically to $1.3633 \cdot 10^5$. During operation, the system maintains successful transmissions for 94.84% of the time, 4.89% of

transmissions have a single error, 0.26% have two consecutive errors, and 0.0138% of all burst errors have three consecutive errors. All the implementations with other values of η result in policies that exploit mixtures of actions and could not be analytically extracted—for the POMDP and H-MDP implementations—and are thus extracted through Monte Carlo simulations, as detailed in the previous section.

5.2 Optimal Policies With Scaled Costs

The portion of time spent in states with n consecutive untimely packets, obtained from the simulations, is shown in **Figure 6** as a function of η , from which we can extract several conclusions. Initially, we notice a sharp decay for the portion of time spent in good states when comparing the values with $\eta = 0.07$ and $\eta = 0.2$. In accord, we notice a sharp increase in all bad states, which is most significant for single errors. This manifests in the optimal policy, as a reluctance of mitigating single burst errors ($n = 1$). Notwithstanding this increase in single errors, all approaches still mitigate higher orders of error bursts ($n > 1$) when $\eta = 0.2$. This is not true for $\eta = 1$, where all approaches focus solely on mitigating the last error that may lead to exceeding the survival time N .



It is important to notice that due to the fact that the H-MDP treats the GE model as hidden when a portion of it is unobservable, turning off an interface results in fully losing the state for that interface. This leads to a behavior where the H-MDP would intentionally turn off the interface that observes a bad state, even when there is no negative incentive to keeping that interface on, in favor of the more likely transition to the steady state of the good state for that interface. Due to this, the optimal policy for the H-MDP is never $\pi^*(S) = A(1, 1) \forall S \in \mathcal{S}$, even when $\eta = 0$. Interestingly, the H-MDP uses the same policy for all three $\eta = 0, 0.03, 0.07$ but is best fit for $\eta = 0.07$. This makes the H-MDP the only suboptimal approach—out of all four—for $\eta = 0$.

Due to the low cost of using the Wi-Fi interface, and generally the superior r probability, all policies keep the Wi-Fi interface at 100% utilization. On the contrary, LTE utilization is the only one that varies for each approach and is shown in **Figure 7**.

However, we are interested in investigating the lifetime of the system given burst error tolerance of N . The system lifetime for the different scaling factors η is given in **Figure 8** as a difference from the optimal system lifetime. Therefore, the goal of each approach with limited observability is to follow the performance of the optimal, fully observable approach as closely as possible. Thus, when a policy improves the system lifetime, it is a sign of energy inefficiency that comes in the form of extra LTE-interface utilization.

Accordingly, the goal of all three approaches that have to work with limited information is to achieve $\Delta K \approx 0$. Looking at **Figure 8**,

we can also notice that, aside from the case of $\eta = 0.07$, the POMDP approach gives the least deviations with regard to the other two approaches. Moreover, the good performance of the H-MDP approach in the case of $\eta = 0.07$ is a simple coincidence since this approach applies exactly the same policy for $0 \leq \eta \leq 0.07$, where both the POMDP and the F-POMDP tend to vary and adapt. Additionally, we notice that the F-POMDP approach is highly focused toward increasing the system lifetime which, as shown in **Figure 7**, comes at the cost of using LTE more often than with the optimal approach. Since this behavior is quite consistent, we can safely say that the F-POMDP is a system-lifetime conservative approach. The POMDP approach is however more adaptable and consistently outperforms the F-POMDP by obtaining a system lifetime that is closely similar to that of the fully observable MDP.

Finally, in **Figure 9**, we show the aggregate deviation, in terms of system lifetime and energy, as calculated as in (28). Here, we can see that the POMDP approach provides the most adaptable behavior, which best resembles the policy when having full observability. Treating the system as a hidden MDP does yield some adaptability; however, the approach can lead to large deviations from the optimal behavior, as it can be seen for $\eta = 0, 0.03$. In these cases, since the stochastic process was treated as hidden to the MDP, the H-MDP optimal solution would intentionally turn off the LTE interface when it is in the bad state. With this, the H-MDP fully loses the information of the LTE interface, in favor of the better stationary-state probabilities. Due to this, we consider the H-MDP approach as




FIGURE 8 | Relative increase in the system lifetime ΔK (i.e., the time to reach one of the absorbing states with N consecutive errors) w.r.t. the fully observable, optimal policy of MDP.


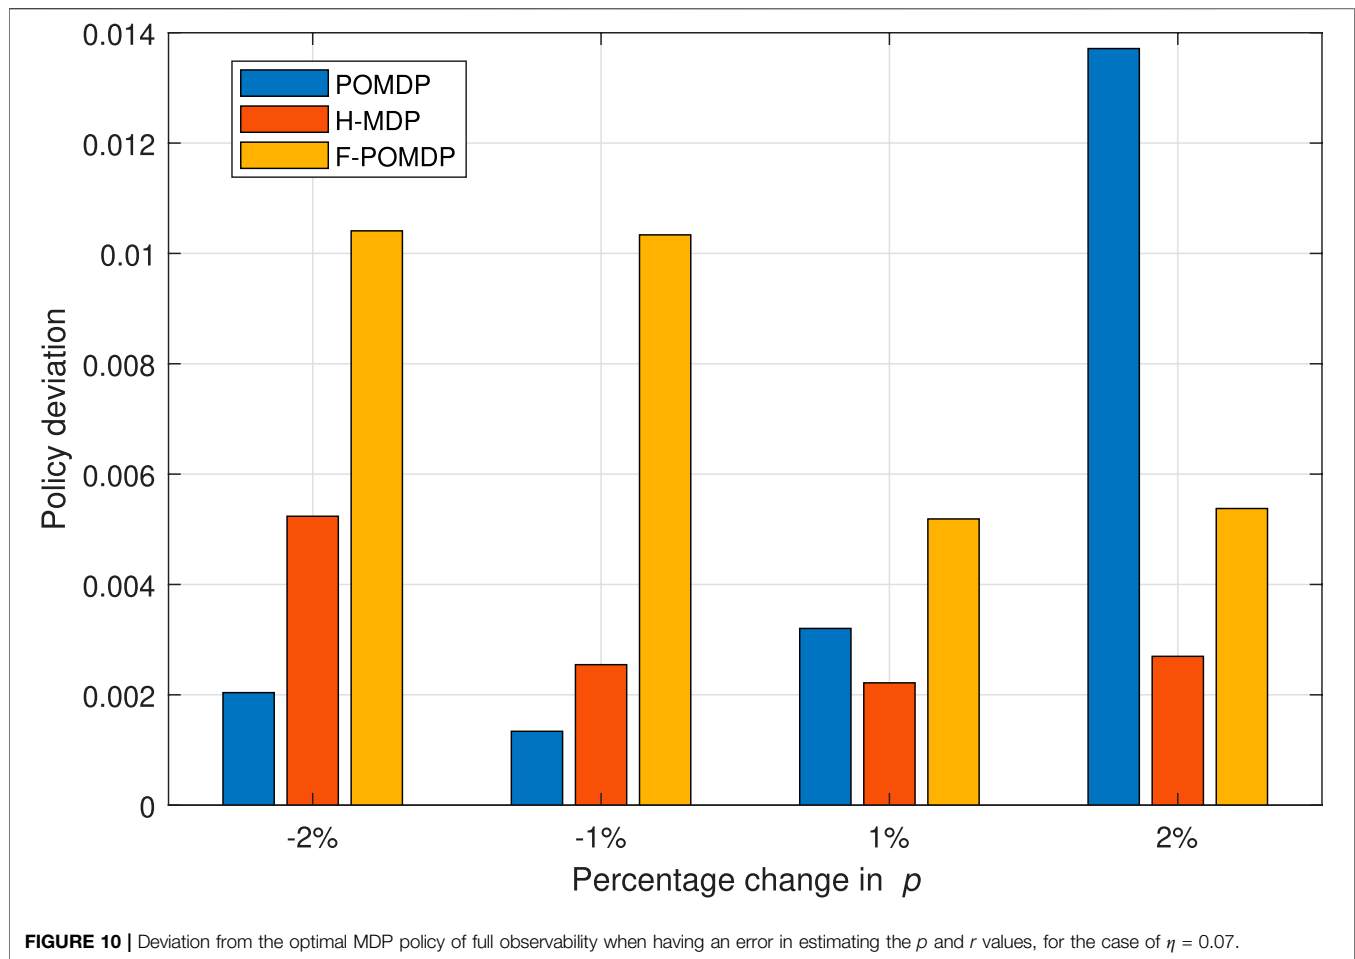


FIGURE 9 | Deviation from the optimal MDP policy of full observability.



unsuitable. On the contrary, the F-POMDP was always conservative with regard to the system lifetime. Thus, the F-POMDP approach is ideal for implementation in devices with extreme power limitations, as it does not require re-computation of the belief states continuously. Finally, the POMDP approach provides the best solutions that most closely follow the optimal policy. Hence, it presents the best solution given that an accurate model of the environment is available and should be adopted if the energy consumption of the computational circuit, which is dedicated for updating the belief states, is not an issue. In the following, we present a sensitivity analysis of the considered methods under an imperfect model of the environment.

5.3 Sensitivity Analysis

We conclude this section by evaluating the impact of the estimation error regarding the p and r values for $\eta = 0.07$. We do this by adding a percentage of error to the p value, while the r value is calculated to maintain the same steady-state probabilities $\pi_{i,B}$ and $\pi_{i,G}$ for each interface i as derived from the values in **Table 1**. In this way, the true probabilities of the Markov system are hidden from the decision processes. As shown in **Figure 10**, in the case of a negative percentage change for the system, the F-POMDP approach greatly deviates from the optimal policy for $\eta = 0.07$. Additionally,

the POMDP approach deviates considerably when an error of 2% is introduced. Moreover, the H-MDP is the most robust as it is more reluctant to change policies in the presence of different parameters, which shows best in the case of positive errors. We can conclude that while belief mechanics help adapt to the optimal policy in the case where the model of the environment is perfectly known, however, such implementations can lead to bad results in particular scenarios where the true probabilities of the system are hidden from the agent. On the contrary, the H-MDP system does not show a big disadvantage in those cases since it already treats the Markov process as hidden. To conclude, we show the relative reward loss $\mathcal{L}(\pi)$ for the same cases of estimation error, as relative to the optimal MDP policy, in **Figure 11**. Here, we observe that even though the F-POMDP method deviates considerably from the optimal MDP policy in the negative estimations of -1-2% (see **Figure 10**), its rewards are close to those with the optimal policy. The reward loss with the H-MDP is relatively stable and does not exceed 0.6%. Finally, we see that the POMDP implementation generally achieves a small reward loss, which is around 0.1% for no error and around 0.4% for -2%. Nevertheless, a high loss is achieved with 2% error. In this case, the adaptability of the POMDP has a negative effect since it scales the policy in accord with the erroneous p and r values.




FIGURE 11 | System rewards with regard to the optimal MDP policy of full observability when having an error in estimating the p and r values, for the case of $\eta = 0.07$.

6 CONCLUSION

Motivated by the recent requirements for cyber-physical systems, we analyzed the problem of addressing error bursts by using two different wireless interfaces. We model the problem as a Gilbert-Elliott model with good and bad states for each interface. Given limited energy resources in our device, we derived and evaluated transmission policies to achieve an adequate trade-off between system lifetime and energy consumption with limited channel information. For this reason, we modeled the system as a POMDP that memorizes and calculates its belief for the observable states. Using value iteration to extract the Q-values from the MDP, we update the policy for the POMDP through the Q_{MDP} technique. Our results show that the POMDP approach indeed produces near-optimal policies when the environment is accurately characterized. As such, this is a computationally inexpensive solution that closely follows the performance of the optimal policy, even in cases with various and mixed state-action pairs. We also propose a forgetful POMDP approach with only two finite belief states. This approach performs worse than the classic POMDP, with affinity to increase system lifetime, but is well suited for approaches that are under extreme energy limitations. Finally, in future works, we would like to practically

validate the usefulness of the system in several application scenarios and address dynamic systems in non-stationary or non-characterized environments.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/supplementary material, and further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

ID involved in formal analysis and software development; ID, IL-M, and JN were responsible for conceptualization, investigation, and writing. ID, IL-M, JN, and PP reviewed and edited the paper. PP and JN involved in obtaining resources, funding acquisition, supervision, and project administration.

FUNDING

This work was supported by the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant, agreement No. 812991, "PAINLESS."

REFERENCES

- 3GPP, 2021 3GPP (2021). Technical specification TS 22.104 V16.5.0. 5G; Service Requirements for Cyber-Physical Control Applications in Vertical Domains. Billinton, R., and Allan, R. N. (1992). *Reliability Evaluation of Engineering Systems*. Springer.
- Chen, Y., Wolf, A., Dörpinghaus, M., Filho, J. C. S. S., Fettweis, G. P., and Oct (2021). Impact of Correlated Fading on Multi-Connectivity. *IEEE Trans. Wireless Commun.* 20, 1011–1022. doi:10.1109/TWC.2020.3030033
- de Sant Ana, P. M., Marchenko, N., Popovski, P., and Soret, B. (2020). “Wireless Control of Autonomous Guided Vehicle Using Reinforcement Learning,” in Proc. IEEE Global Communications Conference (GLOBECOM), Taipei, Taiwan, 7–11 Dec. 2020 (IEEE). doi:10.1109/GLOBECOM42002.2020.9322156
- Dzung, D., Guerraoui, R., Kozhaya, D., and Pignolet, Y.-A. (2015). “To Transmit Now or Not to Transmit Now,” in 2015 IEEE 34th Symposium on Reliable Distributed Systems (SRDS), Montreal, QC, Canada, 28 Sept.–1 Oct. 2015 (IEEE), 246–255. doi:10.1109/SRDS.2015.26
- Haßlinger, G., and Hohlfeld, O. (2008). “The Gilbert-Elliott Model for Packet Loss in Real Time Services on the Internet,” in 14th GI/ITG Conference-Measurement, Modelling and Evaluation of Computer and Communication Systems, Dortmund, Germany, 31 March–2 April 2008 ((VDE)), 1–15.
- Littman, M. L., Cassandra, A. R., and Kaelbling, L. P. (1995). “Learning Policies for Partially Observable Environments: Scaling up,” in *Machine Learning Proceedings 1995*. Editors A. Prieditis and S. Russell (San Francisco (CA): Morgan Kaufmann)), 362–370. doi:10.1016/B978-1-55860-377-6.50052-9
- Mahmood, N. H., Lopez, M., Laselva, D., Pedersen, K., and Berardinelli, G. (2018). “Reliability Oriented Dual Connectivity for URLLC Services in 5G New Radio,” in Proc. International Symposium on Wireless Communication Systems (ISWCS), Lisbon, Portugal, 28–31 Aug. 2018 (IEEE), 1–6. doi:10.1109/ISWCS.2018.8491093
- Nielsen, J. J., Leyva-Mayorga, I., and Popovski, P. (2019). “Reliability and Error Burst Length Analysis of Wireless Multi-Connectivity,” in Proc. International Symposium on Wireless Communication Systems (ISWCS), Oulu, Finland, 27–30 Aug. 2019 (IEEE), 107–111. doi:10.1109/ISWCS.2019.8877248
- Nielsen, J. J., Liu, R., and Popovski, P. (2018). Ultra-reliable Low Latency Communication Using Interface Diversity. *IEEE Trans. Commun.* 66, 1322–1334. doi:10.1109/TCOMM.2017.2771478
- Ploplys, N., Kawka, P., and Alleyne, A. (2004). Closed-loop Control over Wireless Networks. *IEEE Control. Syst.* 24, 58–71. doi:10.1109/mcs.2004.1299533
- Popovski, P., Stefanovic, C., Nielsen, J. J., De Carvalho, E., Angelichinoski, M., Trillingsgaard, K. F., et al. (2019). Wireless Access in Ultra-reliable Low-Latency Communication (URLLC). *IEEE Trans. Commun.* 67, 5783–5801. doi:10.1109/tcomm.2019.2914652
- Puterman, M. L. (2014). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Saad, W., Bennis, M., Chen, M., and Oct (2020). A Vision of 6g Wireless Systems: Applications, Trends, Technologies, and Open Research Problems. *IEEE Netw.* 34, 134–142. doi:10.1109/MNET.001.1900287
- She, C., Chen, Z., Yang, C., Quek, T. Q. S., Li, Y., and Vucetic, B. (2018). Improving Network Availability of Ultra-reliable and Low-Latency Communications with Multi-Connectivity. *IEEE Trans. Commun.* 66, 5482–5496. doi:10.1109/TCOMM.2018.2851244
- Simsek, M., Hößler, T., Jorswieck, E., Klessig, H., and Fettweis, G. (2019). Multiconnectivity in Multicellular, Multiuser Systems: A Matching-Based Approach. *Proc. IEEE* 107, 394–413. doi:10.1109/JPROC.2018.2887265
- Suer, M.-T., Thein, C., Tchouankem, H., and Wolf, L. (2020a). “Evaluation of Multi-Connectivity Schemes for URLLC Traffic over WiFi and LTE,” in Proc. 2020 IEEE Wireless Communications and Networking Conference (WCNC), Seoul, Korea (South), 25–28 May 2020 (IEEE), 1–7. doi:10.1109/WCNC45663.2020.9120829
- Suer, M.-T., Thein, C., Tchouankem, H., and Wolf, L. (2020b). Multi-Connectivity as an Enabler for Reliable Low Latency Communications-An Overview. *IEEE Commun. Surv. Tutorials* 22, 156–169. doi:10.1109/COMST.2019.2949750
- Willig, A., Kubisch, M., Hoene, C., and Wolisz, A. (2002). Measurements of a Wireless Link in an Industrial Environment Using an IEEE 802.11-compliant Physical Layer. *IEEE Trans. Ind. Electron.* 49, 1265–1282. doi:10.1109/TIE.2002.804974
- Wolf, A., Schulz, P., Dorpinghaus, M., Santos Filho, J. C. S., and Fettweis, G. (2019). How Reliable and Capable Is Multi-Connectivity?. *IEEE Trans. Commun.* 67, 1506–1520. doi:10.1109/TCOMM.2018.2873648
- Yajnik, M., Sue Moon, S., Kurose, J., and Towsley, D. (1999). “Measurement and Modelling of the Temporal Dependence in Packet Loss,” in Proc. IEEE INFOCOM ’99. Conference on Computer Communications. Proceedings. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. The Future is Now (Cat. No.99CH36320), New York, NY, USA, 21–25 March 1999 (IEEE), 345–352. doi:10.1109/INFOCOM.1999.749301

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher’s Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors, and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Copyright © 2021 Donevski, Leyva-Mayorga, Nielsen and Popovski. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.