State-Aware Resource Allocation for Wireless Closed-Loop Control Based on Multi-Connectivity

Lucas Scheuvens[1], Philipp Schulz[1], Tom Hößler[2], Norman Franchi[1], André Barreto[2], and Gerhard P. Fettweis[1,2]
[1]Vodafone Chair Mobile Communications Systems, Technische Universität Dresden, Germany
[2]Barkhausen Institut, Dresden, Germany
Email: {lucas.scheuvens, philipp.schulz, norman.franchi, fettweis}@ifn.et.tu-dresden.de
{tom.hoessler, andre.nollbarreto}@barkhauseninstitut.org

***Abstract*—Control communications co-design enables robust and scalable wireless closed-loop control system design. We study the metric "control-communications availability" that allows consecutive errors until the control application is deemed dysfunctional. Multi-connectivity helps increasing the network availability, but there is a lack of dynamic and resource efficient link management. Thus, we propose the "state-aware resource allocation" scheme, whereby parallel links can be assigned adaptively to a given connection, depending on the number of previously, consecutively lost packets. We develop a Markov chain that captures the novel resource allocation approach suited for closed-loop wireless control applications. Our approach outperforms static dual connectivity by two orders of magnitude in terms of control-communications availability while reducing the amount of required resources to approximately half.**

***Index Terms*—Control Communications Co-Design, Wireless Control, Multi-Connectivity, Adaptive Resource Management, Control-Communications Availability**

I. Introduction

Integrating wireless communications into closed-loop control systems is a great challenge for future wireless communications systems and of major significance for future manufacturing, as it will enable a hitherto unknown degree of flexibility and productivity [1]. There currently exist two main approaches in order to achieve this goal.

On the one hand, in the communications research community, ultra reliable low latency communications (URLLC) - as one of the main 5G pillars - is still considered to be the enabling technology, as it is designed to achieve latency values lower than $1\,\mathrm{ms}$ and reliability values larger than $99.9999\,\%$. It targets replacing and extending widely established wired industrial communications systems, such as Ethercat, Profinet I/O, and SERCOS III [2]. In this context, 3GPP defines *reliability* as the complement of the packet loss rate (PLR), i.e., $1-\mathrm{PLR}$ [3], where PLR describes the number of packets that are not transmitted within the time constraint required by the targeted service. However, keeping the PLR lower than 10^{-6} with a latency bound of $1\,\mathrm{ms}$ and one-way payload data rates exceeding $90\,\mathrm{Mbit/s}$ (Ethercat, [4]) will require an immense amount of wireless resources, due to the need for high-bandwidth, high-order multi-connectivity, as well as little to no tolerance for retransmissions (due to the low-latency constraint) [5], [6], still neglecting that retransmissions also cause large signaling overhead. For example, multiple input multiple output (MIMO) channel capacity analysis shows that achieving a capacity of $100\,\mathrm{Mbit/s}$ at a rather large average signal-to-noise ratio (SNR) of $18\,\mathrm{dB}$ with $1\,\mathrm{MHz}$ bandwidth, assuming 100 simultaneous control applications in $100\,\mathrm{MHz}$ system bandwidth, requires approximately 20 antennas at each terminal [7] (still neglecting the aforementioned latency and reliability constraints). This indicates that, especially in a dense industrial setting with hundreds of simultaneous applications, this 1:1 cable replacement approach does not scale.

On the other hand, the control engineering research community has been investigating so-called networked control systems (NCS) for over two decades. These deal with the question of how to cope with communications imperfections on the control side in order to maximize control utility. The imperfections include but are not limited to delay [8], packet drop-out [9], low SNR [10], competition for resources [11], coarse quantization [12], and data rate limitations [13]. A recent survey on results in this domain was conducted in [11]. These works present (sophisticated and advanced) control algorithms for cheap, off-the-shelf, unoptimized communications systems, as the reduced costs for the necessary infrastructure are the main motivator for deploying NCS in the first place.

However, the vision of Industry 4.0 is a high-performing, all-connected factory and it is doubtful that this vision can be put into practice with unmanaged, unoptimized communications systems. At the same time, the scalability issues of URLLC raise the question if URLLC is the technology enabler that will make Industry 4.0 a reality. Hence, we plead for a middle ground between these all-control-domain (NCS) and all-communications-domain (URLLC) approaches. In literature, this is often termed co-design of control application and (wireless) communications (CoCoCo). Specific to our interpretation, it includes the design of wireless communications networks that are able to deal with the trade-off between costly wireless resource utilization and control performance. Hence, we propose a shift away from URLLC towards communications systems that are intercoupled with the control application in order to decide in real-time the importance of each transmitted packet. As we will sketch throughout this article, this has the potential to save a vast amount of valuable resources, depending on the control application under consideration.

This article is structured as follows. In Section II, we introduce the novel key performance indicator *control-communications availability* that describes the (un-)availability of the control application due to (un-)availability of the communications service. In Section III, we present a novel resource allocation scheme that is particularly beneficial for wireless closed-loop control applications because it maximizes the control-communications availability while keeping the average resource consumption low. In Section IV, we develop a Markov chain to model the novel resource allocation scheme before we evaluate the results in Section V. In Section VI, we conclude the article.

II. Control-Communications Availability

The term *availability* commonly denotes the probability of successfully transmitting a packet [14]. It can be increased through the use of frequency, time, space, and/or code diversity. For closed-loop control applications however, diversity in time, i.e., retransmissions of the *same* data – which also extends to hybrid repetitions such as Hybrid Automatic Repeat Request (HARQ) – is often not feasible due to the data being outdated fast, deeming the retransmission useless [8], [15]. Hence, in this work, only diversity schemes supporting simultaneous transmissions are considered and we more specifically limit ourselves to frequency diversity for simplicity.

For a fixed number of simultaneously used, equally distributed and uncorrelated links $L \in \mathbb{N}$, a selection combining (SC) scheme is considered because of its low complexity, which enables the combination of links in higher network layers, e.g., medium access control (MAC). All links are assumed to have a frequency spacing larger than the coherence bandwidth and the packet interarrival time T_s is assumed larger than the coherence time such that all transmissions can be regarded as independent in both, frequency and time. With these assumptions, the availability is given as

$$A_\mathrm{com} = 1 - p_\mathrm{loss}^L \tag{1}$$

with p_loss describing the per-link packet loss probability. This equation describes a communications-centered metric (as denoted by the subscript), and therefore A_com does not carry information about the availability of the control application. However, in the context of industrial automation, the application-centered availability is of particular interest. Fundamentally, two addenda are required for conversion.

(a) As control applications usually manipulate the physical world, they are potentially dangerous when operating incorrectly, so continuing operation in a failed state might have severe consequences. In order to restart the control application after a failure, a “restart procedure” needs to be established that plays a vital role in a meaningful control application availability definition. For modeling purposes, it suffices to define a mean down-time (MDT) that describes a time period for which the control application waits in the “down” state before resuming operation.

(b) It was shown in [16] that the automated guided vehicle (AGV) control application is able to tolerate packet losses as long as not *too many* packets are lost consecutively. Hence, in the following, $K \in \mathbb{N}_0$ will denote the maximum number of consecutive packet losses the control application is able to tolerate. The value of K depends upon the specific control application dynamics and chosen sampling rates and can be derived from a fundamental understanding of the control loop. As a value of $K > 0$ will inevitably degrade control performance, it is left to control engineers to determine meaningful control application requirements and to deduce K for communications system design.

Following the availability definition in [17], availability can also be defined through the mean time to failure (MTTF) and MDT through

$$A_\mathrm{com} = \frac{\mathrm{MTTF}_\mathrm{com}}{\mathrm{MTTF}_\mathrm{com} + \mathrm{MDT}_\mathrm{com}} \tag{2}$$

$$\mathrm{CCA} = \frac{\mathrm{MTTF}}{\mathrm{MTTF} + \mathrm{MDT}} \quad . \tag{3}$$

Thereby, the control-communications availability (CCA) describes the long-term average propability of the *control application* being in the functional “up” state resulting from not loosing more than K consecutive packets. Each time more than K consecutive packet losses occur, the control application is considered “down” for a duration of MDT.

The quantities $\mathrm{MTTF}_\mathrm{com}$ and $\mathrm{MDT}_\mathrm{com}$ (MTTF/MDT with regard to the communications network) for transmissions over L parallel links can be calculated from A_com in a straightforward way via [18]

$$\mathrm{MTTF}_\mathrm{com} = \frac{1}{1 - A_\mathrm{com}} \quad \text{and} \quad \mathrm{MDT}_\mathrm{com} = \frac{1}{A_\mathrm{com}} \quad . \tag{4}$$

For $K = 0$, the MTTF and the $\mathrm{MTTF}_\mathrm{com}$ are identical, while for all other values of K, $\mathrm{MTTF} > \mathrm{MTTF}_\mathrm{com}$. This is because the communications-centered frequency-only diversity (in L) is now complemented by time diversity (in K) *on a control level*, which should not be confused with time diversity on a communications level as introduced at the beginning of this section. We emphasize that this constitutes a try-once-discard (TOD) transmission strategy and *not* retransmissions. As motivated e.g. in [8], [15], at each control instant it is most valuable to transmit the most current instead of outdated data. Quantifying the MTTF for $K > 0$ is a main contribution of this article.

By adjusting the number of parallel links L, both availability metrics, A_com and CCA, can theoretically be chosen arbitrarily close to 1. However, as a second main contribution of this article, instead of trading the additionally gained time diversity (dimension K) for a reduction of parallel links (dimension L) in order to achieve a target CCA value, we propose to dynamically adjust the number of parallel links L depending on the control application’s state. We will show in the following that a dynamic assignment improves the CCA by orders of magnitude.

Table I
SARA EXAMPLE SCHEMES

	scheme	consecutively lost packets k					
		0	1	2	3	4	$\cdots$
number of links L	S_1^0	1	1	1	1	1	$\cdots$
	S_2^0	2	2	2	2	2	$\cdots$
	S_3^0	3	3	3	3	3	$\cdots$
	S_1^1	1	2	3	4	5	$\cdots$
	S_2^1	2	3	4	5	6	$\cdots$
	S_1^2	1	3	5	7	9	$\cdots$
	S_2^2	2	4	6	8	10	$\cdots$

III. STATE-AWARE RESOURCE ALLOCATION

We introduce state-aware resource allocation (SARA), a resource allocation scheme that targets achieving extremely high CCA values at a low long-time average resource consumption.

With the control loop characteristics in mind, we propose that the application-specific importance of a piece of information should play a significant role on how many wireless resources are allocated for transmission. Essentially, a packet can be deemed imperative if its loss will cause an application failure. In this case, many resources should be used for transmission, consequently increasing the likelihood of successful reception. On the other hand, when a packet is considered “nice to have”, i.e., if the control application can still function properly (although degraded) even with this packet lost, only few resources need to be spent.

Thereby, the importance of a packet transmission is evaluated according to the previous consecutive packet losses. For A long sequence of packet losses makes the correct reception of the following packet extremely important, and triggers the assignment of many links in parallel. This increases the likelihood of successful reception, whereas only a single link is assigned whenever the last transmission was successful.

There are also other ways to determine the importance of a packet, e.g., by emulating the control application on a nearby server (model-predictive control) in order to calculate the predicted impact of a lost packet. However, the sketched approach has the advantage that the communications network alone can decide on the importance of a packet since it knows the success/failure of every packet transmission. Although this may not be as effective as model-predictive control, this does not require the high complexity associated with running real-time control models in software and also circumvents the issue of additional delay due to this complexity.

In SARA, different adaptation schemes are denoted as $\mathrm{S}_l^j(K)$, with l indicating the base number of links, i.e., the number of links allocated after a successful transmission; and j indicating the number of links added for each lost packet. S_L^0 corresponds to a multi-connectivity approach with L fixed links, termed *static* schemes in the following. Whenever a packet is transmitted successfully, the number of links is reset to the base value of the scheme. In the rest of this paper $K = 3$ is considered, which denotes a conservative value considering the design guidelines of [19] that recommend a 10- to 20-fold oversampling rate, and the index K is dropped from the notation. As examples in this article, we consider the schemes described in Tab. I. However, the analysis can be performed for any values of l, j or K and also for arbitrary schemes that do not follow the pattern described above, e.g., schemes that ignore the first lost packet and only react after the second or third.

In this article, we present a model for SARA yielding closed-form expressions to describe (a) the unavailability of the control application, (b) the MTTF, (c) the average PLR, and (d) the average number of links for any given application-aware resource allocation scheme. These metrics constitute important key performance indicators (KPIs) in the realm of CoCoCo, as they enable a meaningful assessment of control application failure rate versus the accompanying cost (in terms of number of links). They also demonstrate that under the given assumptions, the PLR is not the most important indicator for an assessment of control application functionality, contrary to prevailing practice.

IV. MODELING SARA

During operation, every control instance features a history of successful/failed transmissions. When all transmission instants are known, the receiver is at all times able to state the sequence of previous packet transmissions which holds binary information about the success/failure of each attempted transmission.

For the introduced approach of this article, only the sequence of previous packet transmissions until the most recent successful transmission will be of significance, which directly translates to the number of previously consecutively lost packets. Depending on the number of tolerable packet losses K of the underlying control application, we define a set of $K + 2$ control application states that we gather in a state vector $\boldsymbol{s} = [s_0 \quad s_1 \quad \cdots \quad s_{K+1}]$. Thereby, all states except s_{K+1} denote states in which the application works within predetermined boundaries (spanning the CCA domain), whereas in s_{K+1} the application is not considered to operate correctly. Whenever a packet is lost, the control application jumps from state s_k to s_{k+1}, whereas when a packet is transmitted successfully, the control application jumps back to s_0. Since our model is memoryless and the sequence that leads to a certain state does not influence the transition probabilities, the discrete-time Markov model in Fig. 1 is considered. The sequence of previous packet transmissions is assumed to be known to the transmitter from acknowledgements sent by the receiver, which might be erroneous and impact the system. In future works, the model must be extended to allow for erroneous ACK/NACK data, however in this article, these transmissions are considered ideal for simplicity.

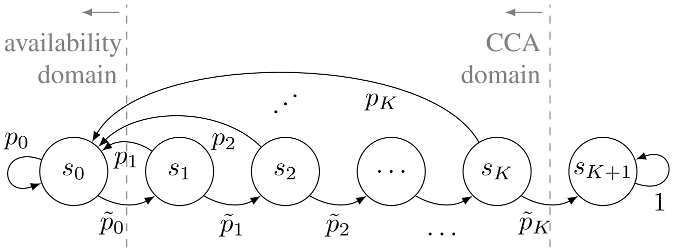

Figure 1. Markov model for improving CCA

The transitions p_k describe the probability of successfully decoding a packet in state s_k. Accordingly, $\tilde{p}_k = 1 - p_k$ denotes the probability of failure. Thereby, p_k equals A_{com} from (1) with L referring to the number of links assigned in s_k. For the adaptive schemes presented in this article, p_k increases for increasing k, while for static resource allocation p_k remains constant.

Let $\boldsymbol{P}$ denote the transition matrix according to Fig. 1, where the row index corresponds to the source state of the transition and the column index to the sink state. All zeros of the matrix are omitted for readability:

$$\boldsymbol{P} = \left[\begin{array}{cccc|c} p_0 & \tilde{p}_0 & & & \\ p_1 & & \tilde{p}_1 & & \\ \vdots & & & \ddots & \\ p_K & & & & \tilde{p}_K \\ \hline & & & & 1 \end{array}\right] = \begin{bmatrix} \boldsymbol{Q} & \boldsymbol{R} \\ \boldsymbol{0} & \boldsymbol{1} \end{bmatrix} \tag{5}$$

This matrix is divided into the sub-matrices $\boldsymbol{Q}$, $\boldsymbol{R}$, $\boldsymbol{0}$, and $\boldsymbol{1}$. Thereby, $\boldsymbol{Q}$ contains the transition probabilities between all transient (non-absorbing) states and $\boldsymbol{R}$ the transition probabilities from transient to absorbing states. $\boldsymbol{0}$ and $\boldsymbol{1}$ constitute matrices of appropriate dimensions of all zeros and ones, respectively. In this case, $\boldsymbol{1}$ above is a scalar, whereas $\boldsymbol{0}$ is a row vector.

Subsequently, the mean sojourn time of each transient state before reaching s_{K+1} can be derived through [20]

$$\boldsymbol{N} = (\boldsymbol{I} - \boldsymbol{Q})^{-1} \tag{6}$$

with $\boldsymbol{I}$ as identity matrix of appropriate dimensions. Therein, each entry N_{ij} describes the expected number of times the process visits state s_j when starting in s_i, before an absorbing state is reached, i.e., the control application fails. Assuming the control application is started in s_0, the first row of $\boldsymbol{N}$ is of particular interest. It can be isolated from $\boldsymbol{N}$ through

$$\boldsymbol{N}_0 = \begin{bmatrix} 1 & 0 & 0 & \dots & 0 \end{bmatrix} \boldsymbol{N} \quad . \tag{7}$$

The sum over all entries in $\boldsymbol{N}_0$ equates to the expected number of transitions that occur in the Markov chain before being absorbed and therefore describes the MTTF_d,

$$\text{MTTF}_\text{d} = \boldsymbol{N}_0 \boldsymbol{1} \quad , \tag{8}$$

with $\boldsymbol{1}$ constituting a column vector of all ones of appropriate dimensions. The subscript emphasizes the unitless/discrete nature of the quantity. It follows that during operation of the control application the state probabilities can be gathered in the row vector

$$\boldsymbol{\pi} = \frac{\boldsymbol{N}_0}{\text{MTTF}_\text{d}} \quad . \tag{9}$$

As mentioned, MTTF_d is unitless (number of discrete time steps) and can be converted to a quantity of time through multiplication with the sampling period T_s of the control application.

$$\text{MTTF} = \text{MTTF}_\text{d} \cdot T_\text{s} \tag{10}$$

The PLR after combining (thus different from p_loss) can be derived from the state probabilities as

$$\text{PLR} = 1 - \pi_0 + \frac{1}{\text{MTTF}_\text{d}} \quad . \tag{11}$$

The intuition is that a successful packet transmission must have occurred whenever the state s_0 is entered. The term $\frac{1}{\text{MTTF}_\text{d}}$ accounts for initializing the Markov chain in s_0.

According to Eq. (3), the MDT needs to be defined in order to evaluate the CCA. In this article, we assume that an AGV will be restarted in s_0 one hour after entering the failed state s_{K+1}. This exemplary scenario is supposed to simulate the rather strict case of manual human intervention that may be required to ensure safe operation. MDT can be interpreted as “punishment” for reaching the failed state. If instead the MDT could be reduced to $30\,\text{s}$, all CCA values of this article would improve by 2 orders of magnitude.

Furthermore, by assigning a cost c_k (here: the number of links) to each state s_k, the average cost of the system can be derived, enabling a direct comparison with the cost of (a) not reacting to packet loss and (b) provisioning high transmission success probabilities all the time. The individual values of c_k are gathered in a cost vector $\boldsymbol{c} = [c_0 \ \cdots \ c_K]$. The scalar product of $\boldsymbol{c}$ and $\boldsymbol{\pi}$ yields the average cost

$$\bar{c} = \boldsymbol{c}\boldsymbol{\pi}^\mathsf{T} \quad . \tag{12}$$

V. Evaluation Results

Tab. II shows the PLR, the control-communication unavailability (the complement of the CCA), the MTTF, and the cost for each of the schemes defined in Section II for a per-link packet loss probability of $p_\text{loss} = 10\,\%$ and a selection combining scheme. Although for all adaptive schemes (S_1^1, S_2^1, S_1^2, S_2^2) the PLR is not significantly lower than for the constant schemes (S_1^0, S_2^0, S_3^0), the control-communications availability is orders of magnitude better. This can be verified in Fig. 2, which shows the control-communications unavailability for the whole range from 0 to 4 tolerable consecutive packet losses. The variation of K is included because the value of $K = 3$, determined in [16] for the AGV application, is not universal and may vary for other control applications. The red curves show the control-communications unavailability for a constant number of links, the blue curves for schemes adding one link per lost packet and the green curves for schemes adding two links per lost packet. The framed numbers denote the average number of utilized links. Note that the constant preset MDT leads to non-linear curves even for schemes with a constant number of links.

Especially for a high tolerance of the control application regarding isolated packet losses, the proposed new adaptive schemes provide a high benefit for the control application. For instance, at three tolerable packet losses, the S_1^1 scheme outperforms the S_2^0 scheme by a factor of 100 in terms of control-communications unavailability (10^{-3} vs. 10^{-5}), while only consuming 1.09 links on average compared with 2 links. This constitutes a shift from MTTF of 34 days to over 10 years for a transmission interval of $T_s = 30\,\text{ms}$ while consuming only approx. half the resources. For two tolerable consecutive packet losses ($K = 2$), the CCA is equal for these schemes but still the advantage in terms of average number of links holds true. Comparing the S_1^2 and S_3^0 schemes at two tolerable packet losses yields a CCA of $1-10^{-4}$ in both schemes, while for S_1^2 only 1.18 links are needed on average compared with constantly 3 links in the S_3^0 case. Also at $K = 2$, comparing S_2^0 and S_2^1, which feature very similar average resource usage (2 links vs. 2.01 links), the improvement in CCA is three orders of magnitude (10^{-1} vs. 10^{-4}), leading to an MTTF increase from approx. 10 hours to roughly 1 year.

Furthermore, we emphasize in Fig. 3 that a high tolerance against consecutive packet losses in the control domain (values of $K \geq 4$) yields extraordinarily high CCA values also for per-link packet loss probabilities exceeding 20 %. The diagram shows the trade-off between CCA and p_{loss} for $K = 2$ and $K = 4$, respectively, across all example schemes of this article. E.g., when targeting a CCA $> 1 - 10^{-5}$ at $K = 4$ (corresponding to a MTTF of approx. 10 years), the SARA scheme S_1^1 enables $p_{\text{loss}} > 20\,\%$. Being able to tolerate on the control application level such high packet loss rates allows to tune the modulation and coding scheme (MCS) on the physical layer of the communications system rather aggressively, enabling a high spectral efficiency per packet transmission that is still able to meet a pre-specified target CCA. However, the details of this physical layer optimization will be left for future work.

VI. Conclusion

This article demonstrates the potential of a new control communications co-design approach that is able to tolerate packet losses to a certain degree on the control side. This, in consequence, significantly relaxes the (wireless) communications requirements regarding PLR. Through SARA, a new, simple yet highly effective and feasible resource allocation approach that exploits this new degree of freedom, the number of simultaneously used links, i.e., the success probability of a transmission, can be adaptively changed, reacting to the success or failure of previous transmissions. This new paradigm enables control-communication availabilities of $> 1 - 10^{-9}$ at a tolerance of $K = 3$ consecutive packet losses at only very little additional resource usage compared to a single-link connection.

Table II
Resulting KPIs for $K = 3$ tolerable packet losses. For MTTF, the sampling period is $T_s = 30\,\text{ms}$.

scheme	average packet loss rate PLR	control-comm. unavailability $1-\text{CCA}$	mean time to failure MTTF	average cost $\bar{c}$
S_1^0	100.0×10^{-3}	9.2×10^{-1}	5 minutes	1.00
S_2^0	10.0×10^{-3}	1.2×10^{-3}	34 days	2.00
S_3^0	1.0×10^{-3}	1.2×10^{-7}	951 years	3.00
S_1^1	91.7×10^{-3}	1.1×10^{-5}	10 years	1.09
S_2^1	9.9×10^{-3}	1.2×10^{-9}	10^5 years	2.01
S_1^2	91.0×10^{-3}	1.1×10^{-11}	10^7 years	1.18
S_2^2	9.9×10^{-3}	1.2×10^{-15}	10^{11} years	2.02

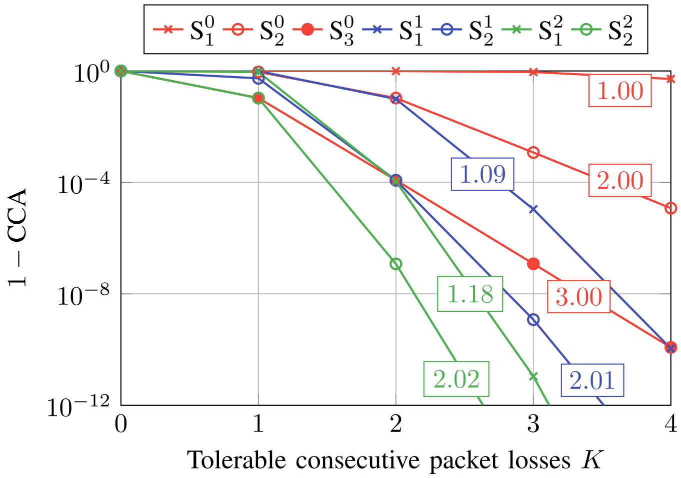

Figure 2. When the control application is able to tolerate K consecutive packet losses, the proposed adaptive increase of links following lost packets reduces the control-communications unavailability by orders of magnitude while the average amount of used links remains low (framed number).

Acknowledgments

This work was in part funded by the German Research Foundation (DFG, Deutsche Forschungsgemeinschaft) as part of Germany’s Excellence Strategy – EXC 2050/1 – Project ID 390696704 – Cluster of Excellence “Centre for Tactile Internet with Human-in-the-Loop” (CeTI) of Technische Universität Dresden. This work was also in part sponsored by the Federal Ministry of Education and Research within the program ”Twenty20 - Partnership for Innovation” under contract 03ZZ0528E - ”fast robotics”. This research was co-financed by public funding of the state of Saxony/Germany.

References

[1] German Federal Ministry of Education and Research, *Industrie 4.0.* Bundesministerium für Bildung und Forschung (BMBF), 2017.

[2] N. Ahmed, A. Mansoor, and R. Ahmad, “Mission-Critical Machine-Type Communications: An Overview and Perspectives towards 5G,” *IEEE Access*, Jan. 2019.

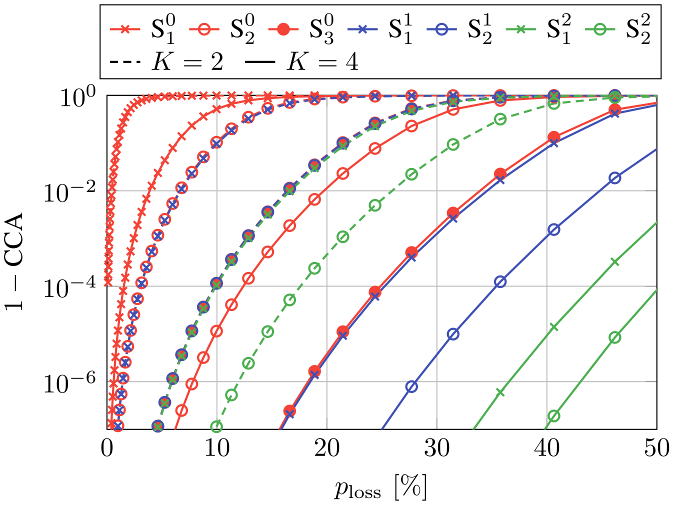

Figure 3. Relationship between p_{loss} and CCA for all example schemes at $K = 2$ and $K = 4$, respectively.

[3] 3rd Generation Partnership Project, “Specification # 22.261 - Service requirements for next generation new services and markets,” 3rd Generation Partnership Project, Technical Specification, 2019, version 16.9.0.

[4] “EtherCAT – The Ethernet Fieldbus,” Beckhoff Automation GmbH, Tech. Rep., Nov. 2018.

[5] R. Abreu, P. Mogensen, and K. I. Pedersen, “Pre-Scheduled Resources for Retransmissions in Ultra-Reliable and Low Latency Communications,” in *Proc. IEEE WCNC*, Mar. 2017.

[6] A. Avranas, M. Kountouris, and P. Ciblat, “Energy-Latency Tradeoff in Ultra-Reliable Low-Latency Communication With Retransmissions,” *IEEE J. Sel. Areas Commun.*, Nov. 2018.

[7] H. Shin and J. H. Lee, “Closed-form formulas for ergodic capacity of MIMO Rayleigh fading channels,” in *IEEE International Conference on Communications*, May 2003.

[8] W. Zhang, M. S. Branicky, and S. M. Phillips, “Stability of networked control systems,” *IEEE Control Syst. Mag.*, Feb. 2001.

[9] J. Hespanha, P. Naghshtabrizi, and Y. Xu, “A Survey of Recent Results in Networked Control Systems,” *Proceedings of the IEEE*, 2007.

[10] J. H. Braslavsky, R. H. Middleton, and J. S. Freudenberg, “Feedback Stabilization Over Signal-to-Noise Ratio Constrained Channels,” *IEEE Trans. Autom. Control*, Aug. 2007.

[11] L. Zhang, H. Gao, and O. Kaynak, “Network-Induced Constraints in Networked Control Systems—A Survey,” *IEEE Trans. Ind. Informat.*, Feb. 2013.

[12] N. Elia and S. Mitter, “Stabilization of linear systems with limited information,” *IEEE Trans. Autom. Control*, Sep. 2001.

[13] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, “Feedback Control Under Data Rate Constraints: An Overview,” *Proceedings of the IEEE*, Jan. 2007.

[14] T. Hößler, P. Schulz, M. Simsek, and G. P. Fettweis, “Mission Availability for Wireless URLLC,” in *Proc. IEEE GLOBECOM*, Dec. 2019.

[15] J. Östman, R. Devassy, G. Durisi, and E. Uysal, “Peak-Age Violation Guarantees for the Transmission of Short Packets over Fading Channels,” in *Proc. IEEE INFOCOM WKSHPS*, Apr. 2019.

[16] L. Scheuvens, T. Hößler, A. Noll Barreto, and G. P. Fettweis, “Wireless Control Communications Co-Design via Application-Adaptive Resource Management,” in *IEEE 5G World Forum*, Dresden, Germany, Oct. 2019.

[17] A. Høyland and M. Rausand, *System Reliability Theory: Models and Statistical Methods*, ser. Wiley Series in Probability and Statistics. Wiley, 2009.

[18] T. Hößler, L. Scheuvens, N. Franchi, M. Simsek, and G. P. Fettweis, “Applying reliability theory for future wireless communication networks,” in *Proc. IEEE PIMRC*, Oct. 2017.

[19] G. F. Franklin, M. L. Workman, and D. Powell, *Digital Control of Dynamic Systems*, 3rd ed. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1997.

[20] C. Grinstead and J. Snell, *Introduction to Probability*. American Mathematical Society, 2012.