

Joint Scheduling and Coding for Reliable, Latency-bounded Transmission over Parallel Wireless Links

Andrea Bedin, *Student Member, IEEE*, Federico Chiariotti, *Member, IEEE*, Andrea Zanella, *Senior Member, IEEE*

Abstract—Several novel industrial applications involve human control of vehicles, cranes, or mobile robots through various high-throughput feedback systems, such as Virtual Reality (VR) and tactile/haptic signals. The near real-time interaction between the system and the operator requires strict latency constraints in packet exchange, which is difficult to guarantee over wireless communication links. In this work, we advocate that packet-level coding and packet scheduling over multiple parallel (unreliable) links have the potential to provide reliable, latency-bounded communication for applications with periodic data generation patterns. However, this goal can be reached only through a careful joint design of such mechanisms, whose interactions can be subtle and difficult to predict. In this paper we first discuss these aspects in general terms, and then present a Markov Decision Process (MDP) model that can be used to find a scheme that optimally exploits the multichannel wireless access in order to maximize the fraction of data blocks delivered within deadline. Our illustrative example is then used to show the optimal coding/scheduling strategies under different combinations of wireless links, also showing that the common solution of backing up a high bitrate unreliable mmWave link with a low bitrate more stable sub-6 GHz link can actually be ineffective in the considered scenario.

Index Terms—Reliability, Multipath Communications, Quality of Service, Tele-operation

I. INTRODUCTION

The new Industry 4.0 paradigm has led to the automation and digitization of many industrial processes, with an increased emphasis on wireless communication technologies: in several industrial and commercial scenarios, the greater flexibility and mobility offered by wireless connections is crucial to perform tasks effectively. The evolution of mobile networks, first with 5G and now with the first steps towards 6G, caters specifically to industrial use cases through the definition of the Ultra-Reliable Low Latency Communications (URLLC) traffic class [1]: the extremely high reliability and predictable latency of cellular wireless access can enable the safe operation of machinery without having to consider communication system faults and imperfections in the control design. However, the stringent constraints of this traffic class require a significant amount of resources, as well as a high level of support

from the network: without scheduling resources in advance and planning transmissions to avoid interference, delivering packets with extremely low latency becomes an arduous task.

In addition, some applications involve mobility and high throughput, making resource allocation even more challenging. A significant example is the remote maneuvering of mobile devices, such as robotic arms, cranes, forklifts, or pickup trucks inside a warehouse, or at the docks, in a mixed environment with human workers [2]. While basic safety mechanisms have to be implemented locally to avoid accidents (e.g., emergency breaking in case of nearby obstacle), macro maneuvering can be left to a remote (human) operator, possibly equipped with a Virtual Reality (VR) visor or, more simply, a screen to visualize the scene around the robot and take appropriate actions, which are then passed to the robot [3].

In this scenario, high-quality video needs to be streamed to the control station with minimal delay to allow for real-time control [4]. For example, the IEEE [5] specifications for high-performance teleoperation require a 20 ms round-trip delay and 4K video resolution, as better explained in Sec. IV. However, the variability of the wireless channel due to nodes' mobility makes static resource allocation unsuitable to guarantee the strict latency requirements of the application, if not by over-provisioning the communication channel, which is neither efficient nor scalable. At the same time, the consequences of late delivery of data may be significant, since the lack of control information may trigger safety mechanisms that stop the robot (also abruptly) to avoid accidents, yielding waste of time and, possibly, mechanical stresses in case of emergency breaking of heavy equipment. Therefore, an unreliable connection can make the whole application unworkable.

The aggregation of multiple wireless links is an interesting option to provide reliability by putting together unreliable components: failures on one link can be compensated by the others, making the connection as a whole stronger and more reliable. This principle has been exploited at the physical and medium access layers, aggregating the transmissions from multiple Base Stations (BSs) or over multiple subcarriers to strengthen the received signal and exploit the independence of wireless channels [6]. However, this has significant signaling and coordination requirements, and can only be performed with full control of the network. It is also possible to aggregate multiple wireless channels *end-to-end*, considering the network as a black box [7]. This approach is significantly more flexible, as it can be applied over any network or combination of networks, aggregating even different operators

Andrea Bedin (andrea.bedin@nokia.com) and Andrea Zanella (zanella@dei.unipd.it) are with the Department of Information Engineering, University of Padova, Italy. Andrea Bedin is also with Nokia Bell Labs, Espoo, Finland. Federico Chiariotti (corresponding author, fchi@es.aau.dk) is with the Department of Electronic Systems, Aalborg University, Denmark. This project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant agreement No. 861222.

and technologies, but still needs careful optimization, as there are several potential issues.

In this work, we present the main issues and solutions for multipath aggregation, aimed at real-time control use cases with high throughput, significant reliability requirements, and strict application-level latency constraints. The challenges presented by these use-cases cannot be solved by purely access network mechanisms, such as URLLC, as performance has to be evaluated end-to-end, not only on the wireless segment. The solution we propose, then, focuses on application data blocks (e.g., video frames), which can be protected using packet-level network coding and transmitted over multiple independent connections. The transmission of each block can be optimized as a Markov Decision Process (MDP), a model which can capture the effect of sequential actions on the future state of the network [8]. We discuss the main issues affecting this type of strategies, as well as the existing support in currently standardized protocols and networks, and present the main trade-offs in a remote-control use case.

The rest of this paper is organized as follows: first, we discuss the requirements and challenges of reliable low-latency multipath transmission in Sec. II. We then present a possible solution to the scheduling and packet-level coding optimization problem in Sec. III, which we evaluate in a mixed mmWave and sub-6 GHz scenario in Sec. IV. Finally, we conclude the paper and present some potential research directions in Sec. V.

II. RELIABLE MULTIPATH TRANSMISSION

We consider a throughput-intensive application generating data blocks at regular intervals. Remote control applications involving rich feedback, such as video or Augmented Reality (AR), generate this type of traffic: frames are recorded or rendered at a constant pace, but might have different sizes, generally much larger than other sensor data in industrial applications.

We then have two potential scenarios, which depend on the nature of the wireless channels being used:

- *Time-varying* channels have a variable capacity, which can fluctuate depending on the nodes' mobility and propagation characteristics of the environment, as well as on the cross-traffic sharing the same wireless resources. Depending on the amount of data transmitted through the wireless link, channel fluctuations might result in higher delivery latency or packet loss rates;
- *On/off* channels can either deliver all the data transmitted through them within the deadline, or be completely unavailable. Links at mmWave frequencies are likely the most significant example of this behavior in modern networks [9], as blockages can make the direct link completely unavailable.

In both cases, a single wireless link cannot guarantee reliable transmission within a tight latency constraint: time-varying channels might only deliver part of the data on time, while on-off channels might drop the entire frame, but the overall effect on the application is the same. Traditional end-to-end protocols such as TCP or QUIC deal with this issue using

Automatic Repeat reQuest (ARQ), i.e., waiting for the packets to be acknowledged and retransmitting lost or excessively delayed [10] data. However, retransmissions have a significant impact on the frame delay, and can cause the Head of Line blocking (HoL) problem: since packets and blocks need to be delivered in the correct order, packet loss can negatively affect even future frames, as it will prevent the application from getting the data until the retransmission is successful.

In these cases, packet-level Forward Error Correction (FEC) is a powerful alternative, exploiting redundant information to allow the receiver to reconstruct missing data. On a single channel, this can mitigate the effects of packet loss, but it is ineffective if the issue is a drop in the link's capacity, since the block transmission with FEC will take even longer, increasing the delay and the risk of violating the latency constraint. Conversely, FEC can help when using multiple parallel links [7]: if the channels are independent, blockages or capacity drops on one link can be compensated by the others, if the redundancy is sufficient.

Fig. 1 shows an example of the effect of multipath coding in the context of a teleoperation task: an industrial robot observes the environment and renders its camera data into a VR feed, which needs to be transmitted to the user. The data block in the example is composed of 4 original (green) packets that are encoded as 7 coded packets: four packets (cyan) are transmitted on a mmWave 5G link and three (orange) on a backup 4G sub-6 link for redundancy and reliability. The reception of any set of 4 packets out of the 7 transmitted will allow the receiver to recover the original frame and display it to the human operator, increasing the probability of meeting the deadline.

However, exploiting path diversity to ensure reliability poses a difficult challenge: since transmitting more packets on a link can increase queuing and, hence, delay for future frames, the amount of redundancy and the splitting of coded packets among the available links need to be optimized to guarantee both high reliability for the current frame and a relatively low impact on future frames.

A. Self-Queuing Delay

Queuing delay is a key issue for time-constrained applications: if the capacity of a link is exceeded, packets pile up, increasing the delay not only for the current frame, but also for future ones. The self-queuing delay issue is well-known in the end-to-end transport literature [10], and led to the design of several low-latency TCP versions.

We can distinguish three cases, which depend on the nature of the application and on the amount of available network support:

- 1) In *underloaded* links, the capacity of the link is much higher than the application bitrate, and queuing is extremely rare. In this case, self-queuing is not an issue, and even robust coding can be supported easily. When available (i.e., not blocked), the mmWave links can be considered as underloaded;
- 2) In *preemptive* links, the link is under high load, but the application can control whether packets in the queue




Fig. 1: Schematic of VR teleoperation over a mixed 4G/5G multipath connection.

are transmitted or discarded. If the previous block has been received correctly, or can be superseded by the next one (which contains more up-to-date information), queuing is not an issue, as older packets are simply dropped from the queue. However, if the older block is still being transmitted, and is required for decoding the new information (as in most video encoding schemes), queuing can become a problem. Furthermore, preemptive operation requires a high level of integration with the network, which might not always be available;

- 3) In *uncontrolled* links, packets that are sent through the link are only discarded in case of buffer overflows, and congestion can have a significant impact on the delay of future frames. This is the case in most public connections, in which the networking infrastructure is not directly controlled by the system designer.

These three cases are increasingly complex to optimize, as actions have deeper consequences on the future viability of the link: while transmitting too many packets through an underloaded link has almost no consequences, doing the same over an uncontrolled link can affect multiple future frames, causing a significant performance drop [11].

B. Link Quality Estimation

In order to effectively determine the amount of redundancy needed to transmit a frame reliably and within the deadline, we need to have an estimate of the distribution of future link capacity: in on/off links, this is equivalent to estimating the blockage probability, as packets are always delivered on time when the link is available. On the other hand, loaded, time-varying links present a tougher challenge: in order to optimize the redundancy, we need an estimate of the Probability Mass Function (PMF) of the number of packets that can be delivered over them in a certain time interval. This might be partly obtained by querying the network interface firmware for the link quality information, but cross-traffic is harder to estimate.

In general, a higher uncertainty on this PMF is particularly damaging for uncontrolled links, as it increases the amount of redundancy necessary for full reliability, inevitably affecting future frames due to self-queuing delay. In addition, network support is less likely for uncontrolled links, so the link quality estimation might have to be performed entirely over the top, with no direct information on the signal quality or state of the buffer.

Over the top estimation mechanisms use ACK packets to infer the available end-to-end connection capacity without any cross-layer input, and are exploited by congestion control algorithms such as Google’s Bottleneck Bandwidth and Round-trip propagation time (BBR) [12]. However, they also present some additional challenges, as the scheduling process and the capacity estimation process are inherently intertwined:

- Delayed feedback might have a significant effect on the estimation of both the state of the queue and the channel capacity: as we are considering individual wireless links, this delay might not be large, but it needs to be taken into account;
- Estimating large capacities requires bursty transmission to saturate the channel, which may lead to higher delays. BBR periodically increases its pace to detect capacity increases, but is relatively slow to adapt to stepwise changes, which can be relatively common in wireless links, e.g., due to changes in the Modulation and Coding Scheme (MCS);
- Low-capacity or high-uncertainty links are typically under-used by the scheduler, as they do not provide any guarantees of timely delivery. However, the quality of the estimate also depends on a minimum number of packets being sent to get accurate statistics. As sending fewer packets increases uncertainty, and increased uncertainty leads schedulers to reduce the number of packets sent through the link, the feedback loop can snowball into making high-uncertainty links entirely unusable.

III. JOINT CODING AND SCHEDULING OPTIMIZATION

We then consider a solution to jointly optimize the coding and scheduling of the blocks. Scheduling more packets on a given link has two effects: the first is that the probability of delivering the block on time increases thanks to the additional redundancy, and the second is that the probability of generating a queue for the next block increases (if the link is uncontrolled). These may not even be at the same time, as the deadline for a block is not necessarily tied to the generation of the next block: VR frames are usually transmitted at very high frame rates, accommodating longer delays in the network, so that multiple frames might be in the transmission buffer at the same time. However, the only moment when the transmitter needs to make a decision is when a new frame is generated, and most applications follow deterministic or at least partially predictable patterns.




Fig. 2: Above: Schematic of VR tele-operation over a mixed 4G/5G multipath connection. Below: example of system model, with two good/bad (Gilbert-Elliot) wireless channels and single-packet blocks.

MDPs are a natural way to model systems making sequential decisions whose consequences can affect future performance in potentially complex ways. The underlying model is a Markov chain, i.e., a stochastic process that takes values in a *discrete state* space and whose future evolution does not depend on the past history. The transition from one state to the next is stochastic and depends on transition probabilities that are affected by the *strategy* adopted by an *agent*, who can observe the current state and take *actions* in a certain set, receiving a *reward* of some kind. The final objective is to find the strategy that maximizes the long-term expected reward. The literature on MDPs is extensive, and there are several methods to solve them and find the optimal *policy* mapping each state to the best possible action in that state: usually, only small MDPs in which the system model is well-known *a priori* can be optimized analytically, and most of the literature focuses on Reinforcement Learning (RL) solutions that learn a policy from experience by trial and error [13].

A natural representation of the reliable multipath communication problem as an MDP is shown in the upper part of Fig. 2, with the following main components:

- The *actions* of the system correspond to the possible packet schedules among the different links: in the graphical example in Fig. 1, encoding the 4 original packets into 7, sending 4 packets over the mmWave link and 3 on the 4G sub-6 one, is a possible action. If the connection includes preemptive links, dropping packets from the queue also needs to be considered as a potential component of the agent’s action set. Similarly, if all links are uncontrolled, dropping a block outright is a potential action that can help ease congestion, at the cost of sacrificing a single block;
- The *reward* function is the on-time delivery of the (de-

coded) block, often communicated through a block-level acknowledgment, regardless of the links the packets were delivered through. As mentioned, most MDP solutions try to maximize the expected reward over long timescales, ensuring that the agent will not act myopically;

- The *state* of the system is the most complex element, as it includes information about the physical state of each wireless channel, the cross-traffic on the link, and the state of the transmission buffer. The state of each wireless channel can possibly be summarized in a single value, such as the Signal to Noise Ratio (SNR) or the Received Signal Strength Indicator (RSSI) or similar, but the uncertainty will increase with the inaccuracy of the information, leading to a lower reliance on the uncertain link. In on/off links, this is simplified to a single value representing the blockage probability.

The transition probability from one state to the other is naturally very complex, and may be impossible to model in realistic systems, leading to the use of RL to find the optimal solution. The lower part of Fig. 2 shows the MDP model of a simple scenario with single-packet data blocks and two parallel binary (good/bad) wireless channels. The system state is hence given by the state of both channels and the action indicates the number of packets to be sent on each channel.

However, the basic model is the same for a complex network with several links of different types, as the only difference is the size of the state and action spaces of the problem. Additionally, we only consider cases in which the state of the link can be estimated directly from the network interface: the case in which estimation is entirely performed end-to-end, as described in Sec. II-B, adds a third effect to the agent’s action, which is to provide an estimate of the state. Systems in which the state is not observed directly and ideally, but its accuracy depends on the agent’s actions, are modeled as Partially Observable MDPs (POMDPs), a more complex extension of MDPs. Strategies also become slightly different: for example, it might be convenient to transmit a few packets over a link known to be in outage, as it is the only way to infer when the outage condition is over and the link becomes usable again [13].

IV. MODEL SETTINGS AND RESULTS

We consider a scenario in which a robot is remotely controlled by a VR operator. The maximum motion-to-photon latency to avoid cybersickness and loss of performance is defined in a recent IEEE standard [5] as 20 ms, which also specifies a required frame rate of at least 120 Frames per Second (FPS). This leaves about 12.5 ms for the transmission of the VR feed, considering realistic delays in the frame generation and feedback transmission. The standard also recommends 4K resolution, which results in a bitrate of approximately 48 Mb/s using the advanced H.265 compression standard [14]. Each frame is then approximately 400 kb, and we can consider encoder settings that result in approximately Constant Bit Rate (CBR) streaming. The main metric we aim at optimizing is the latency-constrained reliability, i.e., the probability that a new VR frame will be delivered within the deadline.




Fig. 3: Latency-constrained reliability as a function of the outage probability of the mmWave links. The capacity of the sub-6 Link is fixed to 36Mbit/s.




Fig. 4: Latency-constrained reliability as a function of the average capacity of the sub-6 GHz links. The mmWave link outage probability is $p_{\text{out}} = 0.2$.

We compare three different cases with two parallel wireless channels, which are assumed to be statistically independent:

- *Pure mmWave*: the robot can transmit the video feed to two BSs in parallel. The channels are considered as on/off, as mmWave has a very high capacity when in line-of-sight, but has frequent outages due to blockages. The channels are then defined by their outage probabilities (p_{out}), since the mean sojourn times in such a state are fixed to 5 steps;
- *Pure sub-6*: the robot can transmit the video over two independent channels in the traditional frequency bands for cellular networks, commonly used for 4G and 5G. These channels are less affected by blockages and other physical obstacles, as lower frequencies have a better penetration, but can easily become congested, as the lower bandwidth can not handle all the cross-traffic. As commonly seen in the literature, the transmission over a sub-6 channel is modeled as an exponentially distributed time, whose parameter is given by the mean link capacity C_{Sub6} divided by the packet size;
- *Mixed*: this is a commonly proposed solution to the problem of mmWave blockage, with a sub-6 channel serving as a backup when there is no line-of-sight mmWave link [15].

The wireless channel models are abstract, as they depend on a single parameter (the outage probability for mmWave links and the average capacity for sub-6 channels). In addition, we assume the transmitter instantaneously knows the conditions of the wireless channels (assuming, e.g., cross-layer communication with the wireless cards or accurate channel-state estimation mechanisms). Such assumptions allowed us to derive the optimal policy and its performance analytically. In a more realistic scenario, this would possibly require more complex solutions such as RL.

While relatively common in the literature, mixed systems using a sub-6 link as a backup for a high-capacity mmWave link actually present a problem: since outages in the mmWave link result in the delivery of no packets, the sub-6 alternative needs to be able to deliver the whole block on its own within the deadline to be an effective backup. Fig. 3 shows the latency-constrained reliability of the systems in the three considered cases as a function of the outage probability p_{out} for

		Link 1 queue				
		0	1	2	3	4
Link 2 queue	0	0.6 0.5	0.5 0.6	0.5 0.7	0.4 0.7	0.4 0.8
	1	0.6 0.5	0.6 0.6	0.5 0.6	0.5 0.7	0.4 0.7
	2	0.7 0.5	0.6 0.5	0.6 0.6	0.5 0.6	0.4 0.7
	3	0.7 0.4	0.7 0.5	0.6 0.5	0.6 0.5	0.5 0.6
	4	0.8 0.4	0.7 0.4	0.7 0.4	0.6 0.5	0.5 0.5

Fig. 5: Strategies for the pure sub-6 GHz system.

the mmWave links. Naturally, having 2 parallel mmWave links works best if the outage probability is low, as the frame can be replicated fully and sent over both. However, if $p_{\text{out}} \geq 0.2$, the two mmWave links will often be blocked at the same time, and the pure sub-6 case, in which both links deliver at least part of the frame and packet-level coding is fully exploited, actually provides a better performance. Interestingly, mixing a mmWave link and a sub-6 link is not the best choice for any value of the outage probability, as another mmWave link would be better if p_{out} is low, and the pure sub-6 scheme can use packet-level coding more efficiently if p_{out} is large.

We can see the same if we fix $p_{\text{out}} = 0.2$ and vary the capacity of the sub-6 links, as we show in Fig. 4: the pure mmWave case has a reliability slightly higher than 95%, but two sub-6 channels can do better as soon as their capacity greater than about 35 Mb/s. The mixed case is only better than the pure sub-6 if the capacity is below 30 Mb/s, with the mmWave link taking most of the responsibility for the reliable delivery.

We can also consider the different strategies for the three scenarios. In the pure mmWave system, the optimal strategy is always to send a complete copy of the frame if the link is not in outage and to avoid sending anything otherwise, to avoid filling the queue. However, even creating a queue is not a problem, as the mmWave links have a significantly higher capacity than what is needed. Fig. 5 and Fig. 6 show the strategies for the other two cases, with $p_{\text{out}} = 0.2$ and $C_{\text{Sub6}} = 36$ Mb/s. The figures show the fraction of the data

		mmWave queue				
		0	1	2	3	4
Sub-6 GHz queue	0	1 0.2	1 0.2	1 0.3	1 0.3	1 0.4
	1	1 0.1	1 0.2	1 0.2	1 0.2	1 0.2
	2	1 0.1	1 0.1	1 0.1	1 0.1	1 0.1
	3	1 0	1 0	1 0	1 0	1 0
	4	1 0	1 0	1 0	1 0	1 0

Fig. 6: Strategies for the mixed system (not in outage).

sent over each link as a function of the number of packets in the two transmission queues: if we use redundant coding, the sum is larger than 1. The cells are also colored based on the amount of redundancy, with darker cells corresponding to a more robust encoding.

As expected, the strategy for the pure sub-6 case, shown in Fig. 5, is symmetrical: since the two links are identical, the optimal scheduling is mirrored if we reverse the states of the queues. The granularity of the optimization was 0.1, i.e., we considered blocks of 40 kb for the encoding, which explains the apparent asymmetry when both queues are empty. In general, the added redundancy is between 10% and 20%, and we can notice that, in highly unbalanced cases, the link with a shorter queue tends to be assigned more packets, allowing the more congested link to transmit fewer packets and flush the queue. We only showed results for queues of up to 4 packets, as other cases occur with extremely low probability: the optimal strategy tries to maintain short queues, sacrificing some reliability to avoid snowball effects in the near future. This is clear from the behavior when 4 packets are in each queue, in which case no redundancy is sent: in more extreme unlucky cases, the frame might even be dropped entirely to allow the links to reduce their queues and improve the situation for future frames.

We can also look at the optimal policy in the mixed case, shown in Fig. 6 when the mmWave link is active: in all cases, all packets are transmitted over this link, while the sub-6 link is only used to transmit additional redundancy. In general, the probability of the system being in one of the states on the right (i.e., of the mmWave link having a significant queue) is extremely low, and these states are only visited after outages. In case of outage, the situation is reversed, as the sub-6 link is used to transmit all packets, while the mmWave link is left unused to avoid building up queues. As we mentioned above, this would be different with an entirely end-to-end capacity estimation: in that case, the transmitter would need to send some probe packets to avoid wasting capacity due to starvation.

In general, the basic principle of the optimal policy is always to use the strongest links first, aiming at avoiding building up queues, but finding the correct amount of redundancy and the correct threshold for shifting the load between links is tricky

even in this simple scenario. However, the optimal strategy also manages to strengthen the reliability of links, using the independence between channels to compensate for delays and losses even when the application has no control over the underlying network.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have presented multipath wireless communications with packet-level coding over multiple links as a possible solution to provide reliable low-latency service to teleoperation and other novel high-throughput industrial applications. We described in detail the main issues and challenges of such a setup, with considerations on different levels of control over the underlying network and for different wireless technologies. We also presented an MDP model to derive the optimal coding and scheduling policy in a generic network, and give a simple but realistic example.

As MDP models are an active topic of research with quick improvements on new RL solutions, the most interesting research direction for communications is the design of the model itself: finding a compact but expressive representation for the state of a link, which might be over different technologies and with different levels of application control over the scheduler and queues, is a significant challenge. The investigation of partially correlated links, such as mmWave with geometric blockers, is also an extremely interesting extension of our idea, which changes the stakes in allocating redundant information.

REFERENCES

- [1] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Network*, vol. 34, no. 3, pp. 134–142, 2019.
- [2] 3GPP, "Service requirements for cyber-physical control applications in vertical domains," Dec. 2021, TS 22.104.
- [3] F. Hu, Y. Deng, H. Zhou, T. H. Jung, C.-B. Chae, and A. H. Aghvami, "A vision of an XR-aided teleoperation system toward 5G/5G," *IEEE Communications Magazine*, vol. 59, no. 1, pp. 34–40, Feb.
- [4] N. Gholipour, S. Parsaeefard, M. R. Javan, N. Mokari, H. Saeedi, and H. Pishro-Nik, "Cloud-based queuing model for tactile internet in next generation of RAN," in *91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE, May 2020.
- [5] "Standard for Head-Mounted Display (HMD)-based Virtual Reality (VR) sickness reduction technology," *IEEE Std. 3079-2020*, Apr. 2021.
- [6] J. J. Nielsen, R. Liu, and P. Popovski, "Ultra-reliable low latency communication using interface diversity," *IEEE Transactions on Communications*, vol. 66, no. 3, pp. 1322–1334, Nov. 2017.
- [7] F. Chiariotti, A. Zanella, S. Kucera, K. Fahmi, and H. Claussen, "The HOP protocol: Reliable latency-bounded end-to-end multipath communication," *IEEE/ACM Transactions on Networking*, vol. 29, no. 5, pp. 2281–2295, Jun. 2021.
- [8] R. A. Howard, *Dynamic programming and Markov processes*. John Wiley, 1960.
- [9] M. Zhang, M. Polese, M. Mezzavilla, J. Zhu, S. Rangan, S. Panwar, and M. Zorzi, "Will TCP work in mmWave 5G cellular networks?" *IEEE Communications Magazine*, vol. 57, no. 1, pp. 65–71, Jan. 2019.
- [10] M. Polese, F. Chiariotti, E. Bonetto, F. Rigotto, A. Zanella, and M. Zorzi, "A survey on recent advances in transport layer protocols," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 4, pp. 3584–3608, Aug. 2019.
- [11] N. B. Shah, K. Lee, and K. Ramchandran, "When do redundant requests reduce latency?" *IEEE Transactions on Communications*, vol. 64, no. 2, pp. 715–722, Dec. 2015.
- [12] E. Atxutegi, F. Liberal, H. K. Haile, K.-J. Grinnemo, A. Brunstrom, and A. Arvidsson, "On the use of TCP BBR in cellular networks," *IEEE Communications Magazine*, vol. 56, pp. 172–179, Mar. 2018.

- [13] X. Xiang, S. Foo, and H. Zang, “Recent advances in Deep Reinforcement Learning applications for solving Partially Observable Markov Decision Processes (POMDP) problems part 2—applications in transportation, industries, communications and networking and more topics,” *MDPI Machine Learning and Knowledge Extraction*, vol. 3, no. 4, pp. 863–878, Jul. 2021.
- [14] P. Akyazi and T. Ebrahimi, “Comparison of compression efficiency between HEVC/H. 265, VP9 and AV1 based on subjective quality assessments,” in *10th International Conference on Quality of Multimedia Experience (QoMEX)*. IEEE, 2018.
- [15] Y. Liu, J. Liu, A. Argyriou, and S. Ci, “MEC-assisted panoramic VR video streaming over millimeter wave mobile networks,” *IEEE Transactions on Multimedia*, vol. 21, no. 5, pp. 1302–1316, Oct. 2018.



Andrea Bedin is an Early Stage Researcher within the MINTS project, currently working at Nokia Bell Labs in Espoo, Finland. He is also enrolled in the Ph.D. program in information engineering at the University of Padova. He obtained his master’s degree in ICT for internet and multimedia engineering in 2020 and his bachelor’s degree in information engineering in 2018 from the University of Padova. His research interests are in low latency ultra-reliable communications for industrial applications.



Federico Chiariotti [S’15-M’19] is currently an Assistant Professor in the Department of Electronic Systems, Aalborg University, Denmark. He received his PhD in Information Engineering in 2019 from the University of Padova, Italy. He has published more than 50 peer-reviewed papers and received the Best Paper Award in 4 conferences. His research focuses on the latency-oriented design of networking protocols, the use of machine learning in network, and on goal-oriented communications.



Andrea Zanella [S’98-M’01-SM’13] is a Full Professor at the Department of Information Engineering (DEI), University of Padova (Italy). He received the Laurea degree in Computer Engineering in 1998 and the PhD in 2001 from the same University. He is one of the coordinators of the SIGnals and NETworking (SIGNET) research lab. His long-established research activities are in the fields of protocol design, optimization, and performance evaluation of wired and wireless networks.