This paper has been accepted for publication in the IEEE Communications Magazine, 2021. DOI: pending

QUIC-EST: a QUIC-Enabled Scheduling and Transmission Scheme to Maximize VoI with Correlated Data Flows

Federico Chiariotti, *Member, IEEE*, Anay Ajit Deshpande, *Student Member, IEEE*, Marco Giordani, *Member, IEEE*, Kostantinos Antonakoglou, Toktam Mahmoodi, *Senior Member, IEEE*, Andrea Zanella, *Senior Member, IEEE*

***Abstract*—Progress in communication technologies has fostered the development of advanced, interactive applications that require multi-sensor data transmission with low latency and high reliability. Since these requirements are not guaranteed by current application-agnostic transport protocols, these applications mostly rely on customized, application-level scheduling and flow control mechanisms, which lack generality and transparency, making it difficult to jointly control the information flows of different applications. In this work, we propose a unified framework to support the transmission of correlated data flows. We assume that the applications are able to describe the correlation among their data streams and the related service requirements in terms of a Value of Information (VoI) matrix. Hence, we propose QUIC-EST, a transmission scheme that combines the congestion control and multi-stream features of the recently proposed QUIC transport protocol with a proper scheduling algorithm to maximize the VoI at the receiver. To illustrate the idea, we propose the analysis of two relevant use cases, namely inter-vehicular and haptic communications, and demonstrate through simulations how the proposed approach can significantly outperform current transport schemes.**

***Index Terms*—QUIC, latency, reliability, multi-sensory, Vehicle-to-Everything (V2X), Tactile Internet (TI).**

I. Introduction

The next generations of cellular networks (5G and beyond) are expected to support new, challenging interactive applications that, besides the freedom of movement given by wireless connectivity, generally require the timely and synchronized delivery of a multitude of sensor data and commands to guarantee interactivity and control effectiveness [1].

For example, haptic communication allows users to interact with remote environments using haptic sensors and actuators that exchange kinesthetic and tactile information. In the case of closed-loop bilateral teleoperation systems, kinesthetic data is time-sensitive. Although stability control mechanisms can be employed to compensate for end-to-end delays that can perturb the stability of such systems, this approach may deteriorate the *transparency* of the service, i.e., the feeling of interactivity and, hence, the quality of telepresence [2]. A more transparent way to decrease the end-to-end delay, instead, consists in reducing the sensor data to be transmitted according to human perception models, but at the cost of a less accurate reconstruction of the signal at the receiver [3]. Somehow similarly, connected vehicles can exchange data generated by on-board sensors via Vehicle-to-Everything (V2X) communications, in order to collaboratively build a richer context awareness and coordinate driving decisions. However, disseminating the sensors' observations is expected to increase data traffic in vehicular networks by multiple orders of magnitude, thus potentially leading to congestion, so that proper data transmission schemes are needed [4]. Other applications that generate streams of correlated data are, e.g., remote control of swarms of mobile robots, tele-monitoring of industrial processes, immersive interactive virtual reality.

In order to operate effectively, these applications need flow control strategies to avoid delays and packet losses caused by congestion, as well as recovery mechanisms to protect particularly valuable data. The most common transmission protocols, namely, the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP), offer complementary services that, however, are not adequate for the purpose. By using TCP, applications can delegate congestion control and packet retransmissions to the transport layer, which provides a simple and well-tested interface with standardized behavior. However, most congestion control mechanisms can create significant latency issues. tThe TCP in-order delivery constraint can cause the head-of-line blocking problem when all the data streams are multiplexed into the same TCP connection. Indeed, if a packet from any of the sensors is lost, successive packets from all other sensors are buffered at the receiver and released to the application only after that the previously lost packet is successfully recovered. Conversely, UDP offers full flexibility to the applications, but leaves the burden of managing congestion and retransmissions to them.

In order to overcome the issues of these protocols, here we propose QUIC-Enabled Scheduling and Transmission (QUIC-EST). QUIC is a recently developed transport protocol that allows data to be sent in parallel and logically independent *streams*, thus avoiding the head-of-line blocking problem among different streams. This reduces unnecessary delays in the reception of the data, particularly when the number of

F. Chiariotti (corresponding author, email: fchi@es.aau.dk) is with the Department of Electronic Systems, Aalborg University, Denmark. A. A. Deshpande, M. Giordani and A. Zanella (email: name.surname@dei.unipd.it) are with the Department of Information Engineering, University of Padova, Italy. K. Antonakoglou and T. Mahmoodi (email:name.surname@kcl.ac.uk) are with the Centre for Telecommunications Research, King's College London, United Kingdom. This project has received funding from the European Union's Horizon 2020 research and innovation program under the Marie Skłodowska-Curie Grant agreement No. 813999 (WINDMILL).

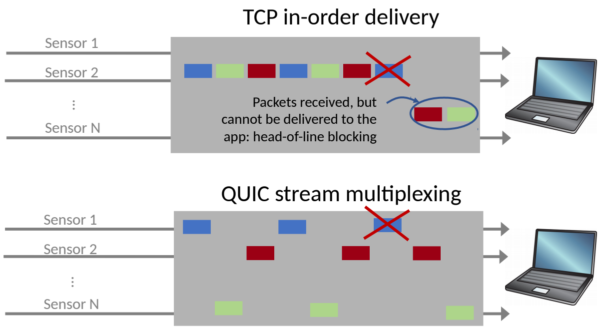

Fig. 1: The head-of-line blocking problem and the stream-based solution.

independent data flows is large. QUIC-EST combines the features of QUIC with a multi-stream scheduling scheme that biases data transmissions as a function of the Value of Information (VoI) provided by the application. Here, the VoI is considered as a scalar value quantifying the utility of the data for the receiver [5]. The VoI takes into consideration the potential correlation of the information flows in time and across different sensors, as well as their intrinsic value. To better illustrate the proposed methodology, we apply it to two relevant use cases, namely autonomous driving and haptic communication, and we show that our approach guarantees better utility compared to traditional transport schemes.

The rest of the paper is organized as follows. Sec. II presents the original QUIC protocol and explains our adaptations. Sec. III describes the proposed VoI-based scheduler. Sec. IV presents the two use cases used to exemplify the potential of the proposed approach. In Sec. V we evaluate the performance of the VoI-based scheduler in such use cases. Finally, Sec. VI concludes the paper.

II. Adapting QUIC for Time-Sensitive Multi-Sensory Applications

The QUIC protocol [6] was designed by Google to solve some of the latency issues that TCP typically causes with Internet traffic. Indeed, TCP offers a single in-order byte-streaming service, leaving the task of separating application-level objects to the application itself. To guarantee in-order delivery, TCP blocks the delivery to the upper layer of any data that has been received out of order, even if logically independent of the lost/delayed packet. QUIC addresses this head-of-line blocking problem by adopting a solution previously implemented by the older Stream Control Transmission Protocol (SCTP) [7], i.e., defining separate streams of data within the same connection. Each stream is treated by the protocol as a logically separate data flow with in-order reliable delivery, independent of the other streams. Fig. 1 shows an example of how QUIC handles multiple streams: while the loss of the blue packet also blocks the orange and green packets in TCP, the logical separation between the streams allows QUIC to deliver the data.

QUIC was designed for Web traffic consisting of a potentially large number of logically independent objects to be delivered with the lowest possible latency. However, its features are also suitable to support interactive multi-sensory applications that need to transmit data from multiple sensors, potentially with low delay, to preserve the user's Quality of Experience (QoE). Nonetheless, unlike Web traffic, sensing data traffic is typically redundant, so that applications do not usually require to receive all the data. This makes the head-of-line blocking issue even more pressing, since the undelivered data might not even be necessary for successful operation. We hence propose the QUIC-EST scheme as a way of adapting QUIC to the multi-sensory application requirements.

In QUIC-EST, each sensory reading can be considered as a separate object. As sensors produce several readings per second, we propose to use not just a different stream for each sensor, but a *different stream for each object*. Whenever all packets sent into a stream are acknowledged, the stream can be reused for a new object. On the contrary, if a stream gets blocked by a packet loss and the data become stale, the sender will transmit a `RESET_STREAM` control frame (which is not bound by in-order delivery constraints) to tell the receiver to discard any existing out-of-order data received from that stream and start again, as suggested in [8].

III. Value of Information-Based Scheduling

While the use of streams allows QUIC to transmit data from different sensors independently, the capacity of the connection might not be sufficient to deliver the data from all sensors within the required time. In this case, the choice of which sensor data should be transmitted and in which order becomes a central problem. Since the QUIC protocol does not specify any scheduler, we propose to implement a priority-based mechanism.

We then define a scheduling algorithm that aims at maximizing the *effective VoI* at the receiver, while avoiding congestion in the connection. To this end, the algorithm needs to be fed with four types of information, namely: *(i)* the (estimated) available capacity of the connection, *(ii)* the size of the data, *(iii)* the intrinsic VoI of the data, and *(iv)* the correlation between the data generated by different sensors (which impacts the effective VoI of the transmitted data). We assume that these input variables are passed to the scheduler using specific interfaces, whose definition is out of the scope of this work. In the following, we provide a more formal description of the variables, and describe their meaning and use.

Let N be the number of objects generated in a batch by an application. Hence, the scheduler is provided with the following inputs.

- The available capacity C along the path, defined as the product of the bottleneck capacity and the minimum round trip time (RTT). These values are estimated directly by the recent bottleneck bandwidth and round-trip propagation time (BBR) congestion control algorithm [9], and can be obtained indirectly when using other latency-aware mechanisms such as the classic Vegas algorithm (which, however, tends to underestimate the capacity in volatile scenarios). In our implementation, we consider the estimate provided by BBR, but note that QUIC natively supports both Vegas and BBR.
- The *size vector* $\mathbf{s} \in \mathbb{N}^N$, which contains the sizes of the objects, in bits. This information is used to check that

the amount of data scheduled for transmission does not exceed the connection capacity C, to avoid congestion.

- The *weight vector* $\mathbf{v} \in [0;1]^N$, which contains the *intrinsic value* of the objects, i.e., the VoI when considering only that source. The intrinsic VoI can depend on a number of factors, such as the position of the sensor (e.g., front sensors in an autonomous vehicle are generally more informative than side sensors for driving decisions, or the finger sensors in a haptic application are more informative than wrist sensors), and the current state of the process (e.g., the presence of an object in a camera's Field of View (FoV), or the detection of a sudden gesture in haptic applications). The intrinsic VoI can also depend on the *time correlation* of the sampling process. If the process is slow-varying, consecutive readings from the same sensor can be highly correlated and, hence, easily predictable by the receiver. Although the relation between the time since the last update from a sensor and the correlation with the new reading is highly application-dependent, it is often assumed to follow an exponential decrease [5]. Some control applications have inbuilt compensation mechanisms for delay, which do not require new measurements until a certain time has passed, so the correlation for these cases can be modeled as a step function. A sigmoid function can then be used to model an imperfect compensation mechanism with a gentler degradation curve. Given the specificities of the different applications, we assume that the intrinsic VoI is determined by the application itself, and passed to the scheduling algorithm in the form of the weight vector. In the next section we will provide examples of how these values can be computed in the two considered use cases.
- The *cross-sensor correlation matrix* $\mathbf{W} \in [0;1]^{N \times N}$, which contains the correlation between objects. Indeed, if the application relies on multiple sensors, there is often a significant amount of redundancy in their information. For example, multiple cameras might have partially overlapping FoVs, or scalar sensors might be measuring correlated quantities. Therefore, the intrinsic VoI of some data may need to be discounted to account for the cross-sensor correlation, because the effective VoI of two correlated measurements can be lower than the sum of the VoIs of the two individual measurements.

The scheduling problem consists in selecting the sensor measurements that provide the maximum VoI at the receiver, while respecting the capacity constraint, i.e., having a total size that is lower than the bandwidth-delay product of the connection. Computing the VoI for all possible scheduling patterns is a combinatorial problem, which soom becomes unfeasible. However, if we limit the analysis to couples of objects, i.e., we do not consider the effects of triplets of correlated objects, this is an instance of the Quadratic Knapsack Problem (QKP) [10], which is NP-hard, but for which there are efficient heuristic solutions.

Fig. 2 shows the basic structure of the proposed scheduling framework: multiple sensors write data with a given VoI to a QUIC socket, and the application supplies the cross-sensor correlation matrix $\mathbf{W}$. The available capacity is read from the BBR estimate, and the scheduler finds the optimal set of objects that can be delivered before the next sensor update, sending them through the connection as fast as congestion control allows. If the connection is lossy or time-varying, scheduling decisions can be revised based on what was already sent, recomputing the solution to the problem.

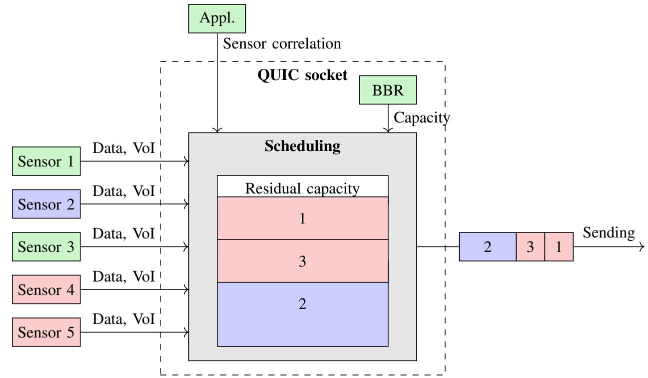

Fig. 2: Basic components of the framework and main data exchanges. In the figure, the data from sensors 1 and 5 is discarded, while the data from sensors 2, 3, and 4 is sent in that order.

To the best of our knowledge, transport layer scheduling of multi-sensory data is an open research problem, which requires the study of the application and sensor features and the estimation of end-to-end capacity dynamics. QUIC-EST forgets the problem by considering correlated measurements in time and across multiple sensors and using congestion control to estimate the path capacity. The scheduling framework is relatively simple, but it can support a wide range of applications, guaranteeing reliability and maximizing the delivered VoI.

IV. Use-case Scenarios for QUIC-EST

The methodology we propose can be applied to any type of application that generates correlated data streams, whose relative importance can be represented in the form of a VoI matrix. In the following, we give two examples of such applications, namely, autonomous driving (Sec. IV-A) and haptic communication (Sec. IV-B).

In our scheduler, data transmission is discriminated based on the VoI provided by the application layer, which depends on the intrinsic characteristics of the different sensors and on the time correlation of consecutive observations. We remark that determining the VoI is not the main focus of this contribution. For the sake of completeness, however, we consider a simple and intuitive definition of the VoI, based on expert knowledge about the relative importance of the different information sources and of their temporal obsolescence. The same empirical approach is used to determine the cross-sensor correlation matrix, which evaluates the degree of correlation among the different sensors. Clearly, more sophisticated strategies may yield different values for the VoI and the correlation matrix that, once fed into QUIC-EST, may result in different transmission strategies. Nonetheless, the values considered in this study are reasonable and apt to illustrate the rationale of the proposed scheme.

A. QUIC-EST for Autonomous Driving

Size vector. The size vector depends on the type of automotive sensor that is considered, the rate at which observations are made, and their resolution. For example, Light Detection and Ranging (LIDAR) sensors can generate data flows with a rate from about 50 kbps to 30 Mbps, depending on the sensor resolution. Similarly, the data rate for camera images ranges from 10 Mbps to 500 Mbps, depending on the image resolution [11], even though compression can reduce the image size by several orders of magnitude. In this work, we consider $N = 5$ sensors: two cameras on the vehicles' top left (lft) and top right (rgt) corners, one on the front (f) and one on the rear side (r), and one LIDAR on the rooftop of the car (L). The sizes of the sensor observations are calculated based on the nuScenes dataset [12], which contains a full autonomous vehicle sensor suite, assuming a 1 Byte pixel encoding and JPG compression for the camera images: we consider a size for the front/rear cameras of 180 KB, for the lateral cameras of 140 KB, and for the LIDAR of 1300 KB, as depicted in Fig. 3 (left).

Intrinsic VoI. We reasonably expect that the LIDAR would be more valuable compared to automotive cameras because it can provide a three-dimensional (rather than bi-dimensional) representation of the environment, and can work efficiently in different weather/time conditions. Also, we assume that the importance of the images taken by the cameras depends on the characteristics of the environment in which the vehicles move (e.g., in the highway scenario lateral cameras will likely make background observations with little marginal information, while frontal/rear cameras might provide more valuable information). Based on these assumptions, we empirically define the correlation vector $\mathbf{v} \in [0,1]^N$ as shown in Fig. 3 (left). Moreover, following the method suggested in [5], we account for the temporal obsolescence of the information by means of an exponential function that depends on the relative age of information, i.e., the time between the generation and reception of the information, with a temporal decay parameter that is proportional to the delay sensitivity of the observation, i.e., the temporal horizon over which that piece of information is considered valuable.

Cross-sensor correlation. We assume that the correlation between images taken by different cameras is proportional to the overlapping of their FoVs. Therefore, the rear camera's images are uncorrelated with those of any other camera. The images taken by lateral cameras are slightly cross-correlated, while higher correlation can be assumed between the frontal and lateral cameras. On the other hand, the LIDAR sensor operates through 360-degree rotations and its observations can be highly correlated/redundant with those of the cameras. The correlation matrix is hence structured as displayed in Fig. 3 (left).

B. QUIC-EST for Haptic Communication

Size vector. In this scenario, the size vector should depend on the number of sensors and actuators integrated on the haptic devices. The CyberGrasp [13] device combines a haptic glove that can sense orientation and movement of the hand and an exoskeleton with five kinesthetic actuators for providing force feedback to the user. Since each haptic glove has 22 movement sensors, considering two hands we have in total $N = 44$ sources of sensor data. Each sensor transmits one floating-point value (i.e., typically 32 bits using the IEEE 754 standard) with a 1 kHz sampling rate, resulting in a 1.4 Mbps total data rate, as represented in Fig. 3 (right).

Intrinsic VoI. In order to determine the VoI of each data sample generated by the haptic device's sensors, we rely on the psycho-physical aspects of human perception. More specifically, we can use Weber's law of Just Noticeable Difference (JND), as in the deadband transmission algorithm in [14], which can be applied in position, velocity and force data. The VoI is then given by the difference between the last transmitted sample from that sensor and the current value, which can be easily computed by the sending application and given to the scheduler. Sensors have the same inherent VoI, but the actual value of the information depends on how novel it is with respect to the one currently available to the receiver. This definition implicitly includes the time correlation between samples, as the difference between consecutive samples will usually be small, but then grow with time consistently with the age of the data.

Cross-sensor correlation. In the haptic communication case, the flexibility of a robotic hand makes the relation between different sensors strongly dependent on their position. If the hand is grasping an object, the correlation between sensors will be different from when it is at rest. Consequently, we cannot give a constant cross-sensor correlation matrix based on the sensors' positions, like we did in the vehicular case. Ideally, the application should be able to compute the instantaneous correlation between sensors in real time and pass the correlation vector to the scheduler. As a simpler (and likely suboptimal) alternative, here we consider the measurements to be independent.

V. Performance Evaluation

In this section, we present a performance evaluation of QUIC-EST, comparing it with other scheduling algorithms in the two scenarios presented in Sec. IV, with extremely different features. While realistic, the assumptions about the two scenarios are arbitrary, and their purpose is to illustrate the methodology from a qualitative perspective, rather than giving a complete quantitative assessment of the scheme. The autonomous vehicle in the first scenario transmits only 10 frames per second but with a maximum rate of 155.2 Mbps; on the other hand, haptic communication has a maximum rate of just 1.4 Mbps, but its sampling frequency is 100 times higher, i.e., 1 kHz. Furthermore, while the haptic communication scenario has 44 different sensors that need to be scheduled, the autonomous driving scenario only has 5.

In both scenarios, we study the average VoI as a function of the available (constant) connection capacity. We consider three other schedulers:

- *First In First Out (FIFO).* This is the default QUIC scheduler, which transmits pieces of data in the same order they were received from the application. It limits

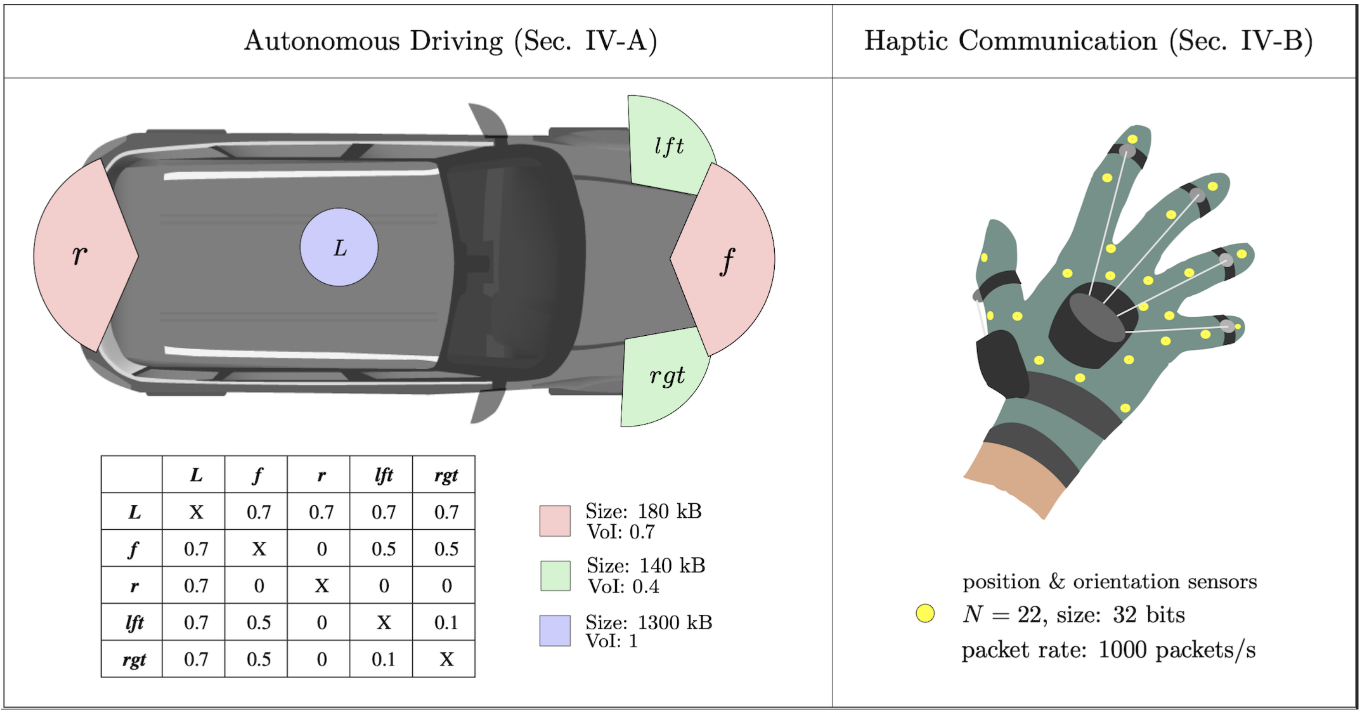

	L	***f***	***r***	***lft***	***rgt***
L	X	0.7	0.7	0.7	0.7
f	0.7	X	0	0.5	0.5
r	0.7	0	X	0	0
lft	0.7	0.5	0	X	0.1
rgt	0.7	0.5	0	0.1	X

Fig. 3: Scheduling input parameters for the autonomous driving (left) and haptic communication (right) scenarios.

transmission to the achievable send rate, discarding any objects that would exceed the connection's capacity. We consider this as a baseline, as its behavior is similar to TCP, without the head-of-line blocking.

- *VoI-based.* This scheduler considers the VoI as the decision factor for transmitting objects that fit the transmission capacity. It is an instance of the classic knapsack problem, as it does not consider cross-sensor correlation or even temporal correlation between values, but only the intrinsic value of each sensor.
- *Cross-sensor VoI.* This scheduler considers cross-sensor correlation, but neglects the temporal correlation. It is equivalent to the optimal scheduler if subsequent measurements from the same sensor are independent, and to the VoI-based scheduler if the measurements are independent between different sensors as well.
- *QUIC-EST.* The scheduler considers VoI as well as time and cross-sensor correlation. The scheduling is obtained by using existing solvers for the QKP. This scheduler gives the best performance if we consider the full application model.

The performance analysis is based on MATLAB simulations, and the code has been made publicly available.[1]

In Fig. 4 we report the normalized VoI achieved by the different schedulers when varying the connection capacities in the autonomous driving scenario. The normalized VoI value is defined as the ratio between the average VoI achieved by each scheduler for a certain channel capacity over the VoI obtained with infinite capacity (which is the same for all the considered schedulers). As expected, the performance grows with the channel capacity for all schedulers, but cross-sensor VoI and

[1]https://github.com/Anay191/Scheduling_Policies_QUIC

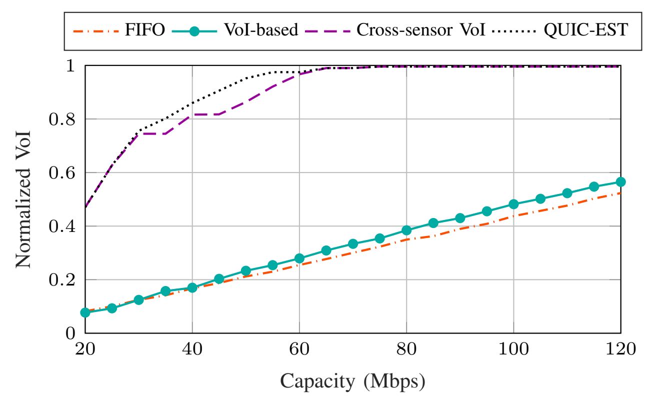

Fig. 4: Normalized VoI for the different schedulers when varying the connection capacity, in the autonomous driving scenario.

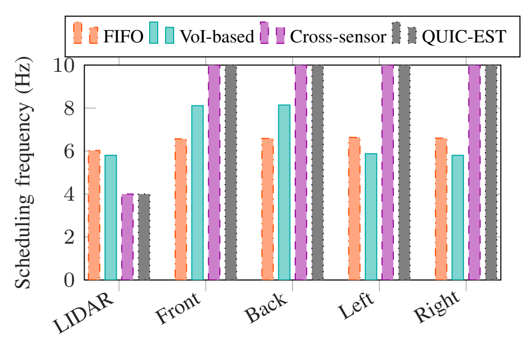

Fig. 5: Average update frequency for the different schedulers in the autonomous driving scenario, with $C = 100$ Mbps.

QUIC-EST, which account for the cross-sensor correlation, exhibit a clear advantage over the others. As Fig. 5 shows, this stark difference is due to the frequency at which the schedulers

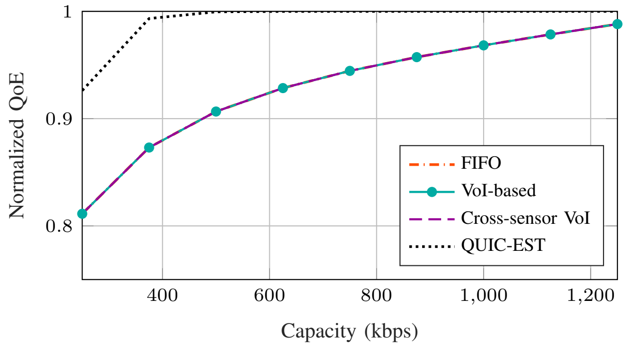

Fig. 6: Comparison between schedulers in the haptic communication scenario in terms of the normalized QoE as a function of capacity.

pick LIDAR frames, which are large and highly correlated with data from the cameras. When the channel capacity is limited, the schedulers that consider cross-correlation among the sensor data flows will limit the number of transmitted LIDAR frames, which are highly correlated, thus leaving more space from camera frames that, with the considered setting, have a higher joint value.

For the haptic communication scenario, we consider the VoI as a logistic function of the difference between the current sample and the last transmitted one. We used realistic haptic traffic model parameters from [15] and cautiously selected a JND value of 5% of the dynamic range of the sensors. Accordingly, we simulate each sensor as an independent Gauss-Markov process, setting $\sigma = 2.15\%$ of the dynamic range to fit the empirical model from the paper. The VoI is then given by a logistic function with center $x_0 = 1.65\sigma$ and sharpness $k = 10$. These values ensure that all sensor measurements that differ for more than the JND are prioritized, while the remaining data are sent only in case of residual capacity. As mentioned, in the haptic communication scenario we neglect the cross-sensor correlation, and all sensors have the same intrinsic VoI, so that the FIFO, VoI-based, and cross-sensor VoI schedulers are all equivalent.

Fig. 6 shows the normalized QoE, defined as the overall fraction of time sensors are under the JND threshold, when varying the channel capacity. We can observe that the QUIC-EST scheme can achieve almost perfect QoE even at less than a third of the capacity needed to send all packets. In this case, the time correlation is critical: the schedulers that do not use a JND-based value, indeed, achieve a lower performance. To be noted that the availability of cross-sensor correlation estimates could further improve the QUIC-EST performance, decreasing the amount of transmission resources needed to support the application.

VI. Conclusions

In this paper, we have presented QUIC-EST, a flexible transmission scheme obtained by combining the emerging QUIC protocol and a VoI-based scheduling strategy and meant for multi-sensory applications with time-sensitive data. We showed that this scheme can be adapted to widely different applications with good results, using autonomous driving and haptic communication as our Future Internet use cases.

At the moment, the design of general transport frameworks for multiple parallel streams of correlated data is still an open research area, whose importance is rising with the popularity of this kind of applications. The combination of such new transport protocols with network slicing techniques is another research avenue of great interest, given the potential to provide applications with reliability guarantees, which can be fundamental for safety-critical services.

References

[1] J. G. Andrews, S. Buzzi, W. Choi, S. V. Hanly, A. Lozano, A. C. Soong, and J. C. Zhang, "What will 5G be?" *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 6, pp. 1065–1082, Jun. 2014.

[2] D. A. Lawrence, "Stability and transparency in bilateral teleoperation," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 5, pp. 624–637, Oct. 1993.

[3] K. Antonakoglou, X. Xu, E. Steinbach, T. Mahmoodi, and M. Dohler, "Toward haptic communications over the 5G Tactile Internet," *IEEE Communications Surveys Tutorials*, vol. 20, no. 4, pp. 3034–3059, Jun. 2018.

[4] T. Higuchi, M. Giordani, A. Zanella, M. Zorzi, and O. Altintas, "Value-anticipating V2V communications for cooperative perception," in *IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2019, pp. 1947–1952.

[5] M. Giordani, A. Zanella, T. Higuchi, O. Altintas, and M. Zorzi, "Investigating Value of Information in Future Vehicular Communications," *IEEE Connected and Automated Vehicles Symposium (CAVS)*, Sep. 2019.

[6] A. Langley, A. Riddoch, A. Wilk *et al.*, "The QUIC transport protocol: Design and internet-scale deployment," in *Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, Aug. 2017.

[7] R. Seggelmann, M. Tüxen, and E. P. Rathgeb, "Stream scheduling considerations for SCTP," in *18th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. IEEE, Sep. 2010, pp. 412–416.

[8] H. Shi, Y. Cui, F. Qian, and Y. Hu, "DTP: Deadline-aware transport protocol," in *3rd Asia-Pacific Workshop on Networking*, Aug. 2019.

[9] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson, "BBR: Congestion-based congestion control," *ACM Queue*, vol. 14, no. 5, pp. 20–53, Oct. 2016.

[10] D. Pisinger, "The quadratic knapsack problem—a survey," *Discrete applied mathematics*, vol. 155, no. 5, pp. 623–648, Mar. 2007.

[11] V. Va, T. Shimizu, G. Bansal, and R. W. Heath, "Millimeter wave vehicular communications: A survey," *Foundations and Trends in Networking*, vol. 10, no. 1, pp. 1–118, Jun. 2016.

[12] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," in *Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE/CVF, Jun. 2020, pp. 11 621–11 631.

[13] M. Aiple and A. Schiele, "Pushing the limits of the CyberGrasp™ for haptic rendering," in *International Conference on Robotics and Automation*. IEEE, May 2013, pp. 3541–3546.

[14] P. Hinterseer, E. Steinbach, S. Hirche, and M. Buss, "A novel, psychophysically motivated transmission approach for haptic data streams in telepresence and teleaction systems," in *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 2. IEEE, Mar. 2005, pp. 1097–1100.

[15] M. Abu-Tair and A. Marshall, "An empirical model for multi-contact point haptic network traffic," in *2nd International Conference on Immersive Telecommunications*, May 2009.

Federico Chiariotti [SM'15-M'19] received his Ph.D. in Information Engineering in 2019 from the University of Padova, Italy. He is currently a postdoctoral researcher at Aalborg University, Denmark. His work received the Best Paper Award in 4 conferences, including the 2020 IEEE INFOCOM WCNEE workshop. His research focuses on the latency-oriented design of networking protocols.

Anay Ajit Deshpande [SM'19] received his Bachelors in Electronics and Communication Engineering in 2016 from VIT University, India and Masters in Information and Communication Engineering in 2019 from Technische Universität Darmstadt, Germany. He is currently a PhD student and a Marie Curie Fellow under H2020-MSCA-ITN-WindMill at University of Padova, Italy. His research focusses on anticipatory techniques in wireless networks.

Marco Giordani [M'20] received his Ph.D. in Information Engineering in 2020 from the University of Padova, Italy, where he is now a postdoctoral researcher and adjunct professor. In 2018 he received the "Daniel E. Noble Fellowship Award" from the IEEE Vehicular Technology Society. His research focuses on protocol design for 5G/6G mmWave cellular and vehicular networks.

Konstantinos Antonakoglou received the M.Sc. degree in Control and Computing from the National and Kapodistrian University of Athens in 2014 and the Ph.D. degree in Telecommunications from King's College London in 2020. He is currently a Research Associate at King's College London and his research interests include haptic communication systems over networks, mobile cloud computing and edge-assisted AI systems.

Andrea Zanella [S'98-M'01-SM'13] is Full Professor at the Department of Information Engineering (DEI), University of Padova (ITALY). He received the Laurea degree in Computer Engineering in 1998 from the same University and the PhD in 2001. 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. .

Toktam Mahmoodi received the B.Sc. degree in electrical engineering from the Sharif University of Technology, Iran, in 2002, and the Ph.D. degree in telecommunications from King's College London, U.K, in 2009. Toktam is currently Head of the Centre for Telecommunications Research at the Department of Engineering, King's College London. Her research interests include mobile communications, network intelligence, and low latency networking.