Towards 6DoF Virtual Reality Video Streaming: Status and Challenges

Jeroen van der Hooft[1], *Maria Torres Vega*[1], *Tim Wauters*[1], *Hemanth Kumar Ravuri*[1], *Christian Timmerer*[2], *Hermann Hellwagner*[2], *Filip De Turck*[1]
[1]*IDLab, Department of Information Technology, Ghent University - imec*
[2]*MMC, Institute of Information Technology, Alpen-Adria-Universität Klagenfurt*
jeroen.vanderhooft@ugent.be

1. Introduction

In the last few years, delivery of immersive video with six degrees of freedom (6DoF) has become an important topic for content providers. Recent technological advancements have resulted in affordable head-mounted displays, allowing a broad range of users to enjoy virtual reality (VR) content. Service providers such as Facebook[1] and YouTube[2] were among the first to provide 360° video, using the principle of HTTP adaptive streaming (HAS) to deliver the content to the end user. In HAS, the content is encoded using several quality representations, temporally segmented into chunks of one to ten seconds and stored on one or multiple servers within a content delivery network. Based on the perceived network conditions, the device characteristics, and the user's preferences, the client can then decide on the quality of each of these segments [1]. Having the ability to adapt the video quality, this approach actively avoids buffer starvation, and therefore results in smoother playback of the requested content and a higher Quality of Experience (QoE) for the end user [2]. The introduction of 360° video provides the user with three degrees of freedom to move within an immersive world, allowing changes in the yaw, roll, and pitch. In the last few years, multiple solutions have been proposed to efficiently deliver VR content through HAS, focusing, for instance, on foveas- and tile-based encoding, improved viewport prediction (i.e., prediction of the user's head movement in the near future, in order to buffer useful, high-quality content), and application layer optimizations [3]. In these works, however, the location of the user remains fixed to the position of the camera within the scene. Recently, significant research efforts have been made to realize 6DoF for streamed video content; i.e., the user may experience three additional degrees of freedom by being able to change the viewing position in a video scene. First progress is promising, but significant research contributions will be required in order to realize its full potential. In this paper, an overview of existing 6DoF solutions is given, and key challenges and opportunities are highlighted.

2. 6DoF Video Solutions

Two types of approaches are generally considered: image-based and volumetric media-based solutions. The former requires a representation of images at every different angle and tilt, while the latter stores objects as a collection of points in the three-dimensional space.

2.1. Image-Based Solutions

In image-based solutions, the system renders different views of an environment from a set of pre-acquired imagery. Images are typically captured using camera arrays (see Figure 1) or cylindrical camera setups, resulting in a representation of images at every different angle and tilt. Because different representations are immediately available, displaying content corresponding to a given position and viewing direction requires modest computational resources [4]. However, the approach results in large storage and bandwidth requirements, since roughly every 0.3-degree difference in angle requires a new image in order to provide a smooth transition between images [5].

While the concept of light fields has been around for two decades, it recently caught more attention as a means to provide 6DoF content streaming; in the last few years, several image-based solutions and frameworks have been proposed. Wijnants et al. propose a DASH-compliant framework for the delivery of light fields [6]. The authors achieve real-time rendering by leveraging video decoding to contemporary consumer-grade GPUs, using disk-versus-GPU caching in order to render source images more quickly. While the proposed framework allows the user to move around and download content in an adaptive manner, only static light fields are considered; delivery of video content is not supported. Furthermore, only single objects are studied in their work. Daniel et al. propose SMFoLD, an open streaming media standard for light field video [7]. This standard allows compliant displays to

[1] https://www.facebook.com, last accessed: August 19, 2019
[2] https://www.youtube.com, last accessed: August 19, 2019

receive a stream of three-dimensional frame descriptions and render scenes without the need for specialized head-mounted devices. The authors discuss several technical challenges to realize a fully functional system framework, focusing on encoding, streaming and displaying the three-dimensional content. Kara et al. propose a framework for subjective evaluation of light field scenes, considering different spatial resolutions and angular differences between images [8]. Similar to traditional video, the authors show that quality switching is preferred over long stalling events. Furthermore, they conclude that it is beneficial to choose a video representation with both spatial and angular resolution reduced, if the smoothness of the continuous horizontal motion parallax cannot be guaranteed. While these research efforts offer interesting insights into the possibilities of light fields for 6DoF content streaming, an operational framework for high-quality image-based 6DoF video streaming does not yet exist.

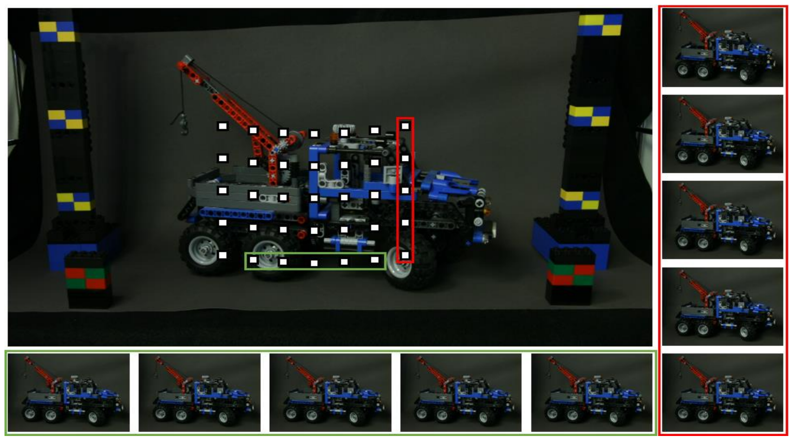

Figure 1. Illustration of light field capture from the (New) Stanford Light Field Archive [9].

2.2. Volumetric Media-Based Solutions

Volumetric media-based solutions store objects as a collection of points. Capturing the geometry (*x, y, z* tuples) and color (*RGB* values) of thousands or millions of points, the object can be rendered from any viewing angle [10]. This reduces storage and bandwidth costs, but requires complex preprocessing (i.e., multiple camera angles and depths) and rendering at client-side. An example capture setup and generated point cloud object are shown in Figure 2.

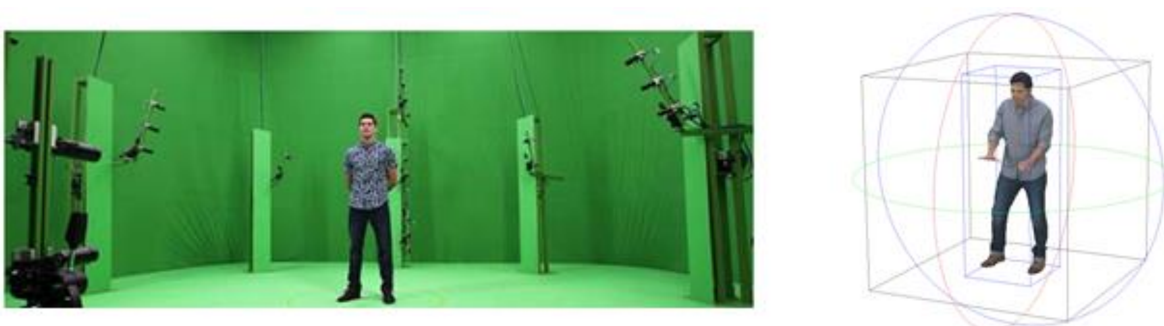

Figure 2. Example point cloud capture from the reference MPEG dataset [11].

Focusing on video on demand, Hosseini and Timmerer are the first to propose a DASH-compliant approach for single point cloud streaming [12]. Rather than using a dedicated encoder, the authors sample the different points to generate versions of lower quality. Furthermore, objects are requested on a per-frame basis, which means that a number of HTTP GET requests proportional to the frame rate is required. He et al. consider view-dependent streaming of point cloud objects, using cubic projection to create six two-dimensional images which can then be compressed using traditional compression techniques [13]. The proposed approach relies on a hybrid network (broadband and broadcast) and in-network optimizations such as caching. In previous work [14]. we propose PCC-DASH, a framework for streaming 6DoF scenes consisting of multiple point cloud objects. Point cloud compression is used to prepare multiple quality versions of the considered objects, and several rate adaptation heuristics are proposed which take into account the user's position and viewing angle. However, the framework is not able to decode and render the content in real-time.

Figure 3. Experimental setup by Dijkstra-Soudarissanane et al. [15]. Two users are immersed in a virtual world, in which they can both see and hear one another in real-time.

Concerning live streaming, Dijkstra-Soudarissanane et al. propose a multi-view end-to-end system for real-time capture, transmission and rendering of volumetric media [15]. This system relies on a multi-point control unit (MCU), to shift processing from end devices into a server. Through a relevant demonstrator, the authors show that two end users can effectively communicate within an immersive world (see Figure 3). Similarly, Qian et al. propose Nebula, a volumetric video system that leverages edge computing to reduce computational efforts on the user's device [16]. The setup uses regular pixel-based video encoding and decoding, in order to stream and render a single point cloud object at the client side. Both solutions show that offloading the decoding to the edge results in timely decoding of smaller point cloud objects, which is of major importance for future volumetric media applications.

3. Challenges and Opportunities

Early 6DoF solutions have shown promising results, but are limited in terms of interactivity and complexity. Below, we discuss different challenges that need to be addressed to fulfill the true potential of 6DoF video streaming. As shown in Figure 4, these challenges focus on i) content representation and encoding, ii) rate adaptation algorithms, iii) application layer protocols, iv) in-network optimizations and v) enhanced evaluation techniques.

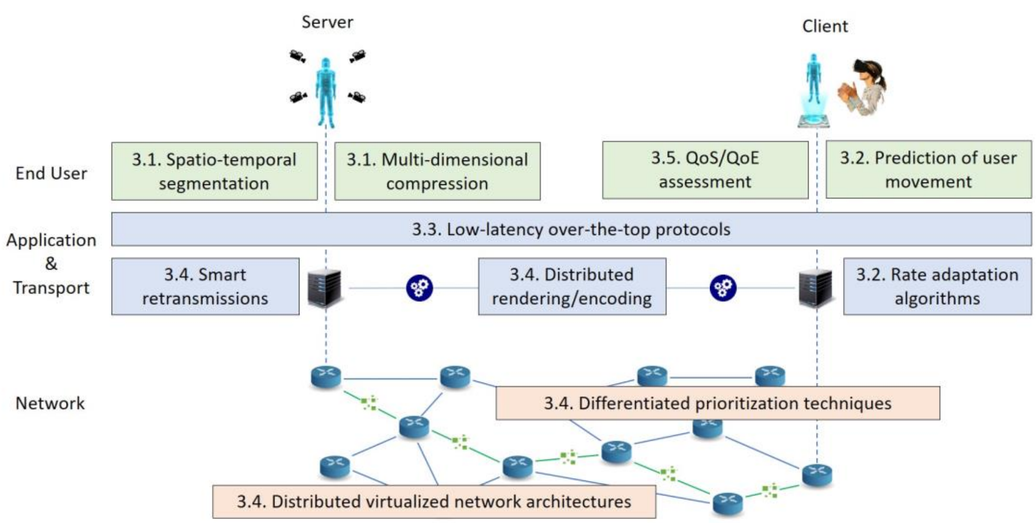

Figure 4. Challenges faced to enable 6DoF video streaming (adapted from [5]).

3.1. Content Representation and Encoding

A key aspect in enabling 6DoF video streaming, is the representation of the considered content. Since 2016, the MPEG project on "Coded representation of immersive media", or MPEG-I, has taken on the challenge of enabling coding, transmission and presentation of immersive media [17]. Efforts have already resulted in the Omnidirectional MediA Format (OMAF), which is currently being extended to support interactivity and 3DoF+, which enables limited modifications of the viewing position. For 6DoF solutions, however, no new formats have been proposed at the time of writing. Therefore, traditional image- and volumetric media-based solutions are still widely adopted.

As mentioned earlier, image-based solutions require a representation of images at every different angle and tilt. Feasible compression rates can be obtained through spatial reduction (i.e., reducing the resolution of each image) and angular reduction (i.e., reducing the number of images, mostly affecting users with high movement) [8]. Meanwhile, compression techniques for point clouds mostly include static kd-tree- and octree-based solutions, with notable examples including Google's Draco [18] and the work by Schnabel and Klein [19]. However, since MPEG launched its call for proposals in 2017 [20], alternative approaches have been suggested. Following an extensive evaluation of nine submitted proposals, MPEG selected a reference encoder for video-based point cloud compression (V-PCC) [10]. This encoder converts point clouds into two separate video sequences, which capture the geometry and texture information, and applies traditional video coding techniques to compress the data. However, this compression technique cannot be used to decode point cloud objects in real-time on commodity hardware [14].

Currently, most encoding techniques are based on compression standards for traditional images and video. Future research efforts should focus on media formats adapted to the considered use case, for instance through the application of sparse representations [21]. Furthermore, as illustrated by the work by Dijkstra-Soudarissanane et al. [15], offloading the decoding to edge nodes offers the advantage of timely decoding and sharing of resources. Ongoing efforts are needed to optimize this process, possibly combining it with in-network optimizations.

3.2. Rate Adaptation Algorithms

Similar to traditional video streaming, rate adaptation is an important factor when enabling 6DoF video streaming over the best-effort Internet. Its complexity, however, is significantly higher; rate adaptation algorithms need to take into account both content characteristics and network throughput and latency, as well as the user’s movement and viewing angle. In this regard, culling the considered objects (i.e., reducing the amount of content based on what regions of the video the user can observe [22]) is a well-known method to reduce bandwidth requirements.

Recently, a number of rate adaptation heuristics have been proposed. Qian et al. present two rate adaptation mechanisms for Nebula [16], while Hosseini presents a rate adaptation heuristic for multiple point cloud objects [23]. There is no focus on evaluation, however, so that no results on the video quality are reported. Park et al. propose a utility-based rate adaptation heuristic for volumetric media, which is both throughput- and buffer-aware [24]. The heuristic was evaluated using simulation, reporting considered utility metric values of each object rather than the resulting visual quality. In our previous work, we propose several rate adaptation heuristics for multi-object point cloud scenes, which consider the user’s position and viewing angle within the scene [14]. Results are encouraging, but indicate that the performance of the heuristics strongly depends on the considered video and camera path.

The above heuristics are typically evaluated using fixed user trajectories. This is an idealized scenario, since the client can adapt the quality of the content based on perfect knowledge. Ongoing research will rather have to focus on predicting the user’s position and viewing direction, based on the user’s (recent) history, video saliency and content. Similar techniques for 360° video already exist [3], but have to be improved to deal with additional complexity.

3.3. Application Layer Optimizations

Nowadays, most HAS solutions use HTTP/1.1 over TCP to retrieve the required resources through request-response transactions, buffering fetched video segments and playing them out in linear order. One possibility to speed up TCP-based solutions is to develop smarter retransmission schemes. In 6DoF video streaming, retransmission is required when transmitting or updating the manifest file or other data crucial for rendering virtual views. When less important data is sent (e.g., incremental data or less important frames), however, the loss may be acceptable. The transport layer needs to handle these different scenarios and decide what to do for which packet, stream or flow.

Some solutions revert to UDP, in order to improve latency at the cost of reliability. One example is the HTTP/3 protocol, which will soon be standardized by the Internet Engineering Task Force (IETF) [25]. This protocol is based on the QUIC protocol, proposed by Google in 2012 [26]. HTTP/3 establishes a number of multiplexed UDP connections, resulting in independent delivery of multiple streams of data. In contrast to HTTP/2, which uses a single TCP connection, this approach avoids head-of-line-blocking if any of the TCP packets are delayed or lost. Another example is WebRTC, a real-time communication protocol which has shown positive results for traditional video in the recent past [27]. WebRTC is, however, peer-to-peer in nature and thus requires multiple encoders at different qualities for each peering connection to ensure an adaptive streaming solution, which hampers scalability. Research on dynamically recomputing encoding settings is very limited and immersive scenarios (with three, let alone six degrees of freedom) have not yet been studied at all.

3.4. In-Network Optimizations

Although various optimizations on higher layers provide support in managing high-bandwidth and low-latency requirements, networks still have a substantial role to play in the end-to-end streaming of 6DoF content. With the advent of multi-camera systems, the computational complexity of tasks such as encoding and rendering increased exponentially, making it difficult to implement them on the end-user equipment. Such tasks can be migrated to resourceful cloud/fog servers on the network. For this reason, networks should be more than just transport circuits.

To support efficient delivery of immersive media services, networks require more programmability. This can be achieved by three evolving technologies: i) software-defined networks (SDN), ii) network function virtualization (NFV) and iii) multi-access edge computing (MEC). The SDN paradigm offers flexibility to the networks in the form of programmable network management, easy reconfiguration and on-demand resource allocation. It suffers, however, from issues such as scalability, control plane overhead and denial of service (DoS) attacks. Such shortcomings are to be addressed in order to support 6DoF content streaming. NFV allows the network functions to be deployed as virtualized software entities running on commodity hardware [28]. Various services involved in 6DoF video streaming can be mapped to respective network functions and can be deployed as a service function chain (SFC). The SFC can be distributed to different locations in accordance to various requirements such as hardware capacity, bandwidth, distance, latency, reliability and their respective tradeoffs [28] [29]. As an example, tasks such as view-synthesis can be offloaded from the end-user equipment but should be placed closer to the user in order to lower the latency. This can be achieved by MEC, which enables the devices to access cloud/fog resources in an on-demand fashion [29]. Since 6DoF video streaming depends on diverse factors, the function placement should be treated as multi-objective optimization problem that has not been extensively researched yet. In addition, there is a need for novel and proactive resource management mechanisms for better resource utilization.

The success of content delivery networks increased the prominence of strategic content caching at the edge. Such storage will play an important role in 6DoF content streaming; upon a new task request the server/network needs to swiftly decide if it should store the content for future requests or not. Furthermore, cache placement and distribution is an important research direction. Proactive caching strategies need to evolve, however, as they depend on spatio-temporal traffic predictions, the users' location, mobility, etc. Other network level approaches such as network coding [30] and network slicing [28] can be exploited to meet the requirements of 6DoF video streaming.

3.5. Evaluation Metrics

The perceived quality of 6DoF video depends on many parameters, such as the frame rate and the degree resolution. Therefore, understanding the effects and "sweet spots" of each of the parameters on the user perception is fundamental to help improve the bandwidth consumption while maintaining the user's QoE [5]. Given the subjective essence of the user's experience, in the last years several works have appeared that subjectively assess the QoE of holographic media. Such are the cases of Kara et al. for adaptive streaming of light field video [8] and of Javaheri et al. [31] for point cloud streaming. However, coping with the dynamics of 6DoF video streaming will require real-time measurements of how the user perceives the streaming. In such cases, objective metrics are better suited for the assessment [31]. For this reason, current effort goes in the direction of devising objective metrics for holographic media correlating with the user's QoE. One first step has been to adapt objective metrics traditionally used for two-dimensional videos for virtual reality and 6DoF video content. Variants derived from the mean square error (MSE) and the peak signal-to-noise ratio (PSNR) have shown good correlation both for point cloud [31] and light field video [32]. Despite the promising results, these studies have focused mainly on encoding derived artifacts. Thus, their accuracy to assess the effects of end-to-end system (capturing, encoding and streaming) degradation is still unknown. Furthermore, most of the current studies (both objective and subjective) assume the user to be passive, thus the effects of the interactivity with the holographic content are not taken into account. This will be a fundamental subject of research for future 6DoF applications.

4. Conclusions

This letter presented a brief overview of ongoing research efforts to realize virtual reality video streaming with six degrees of freedom. Both image- and volumetric media-based representation techniques were discussed, and a list of challenges and opportunities for future work was presented. Given the many applications of virtual reality and an increased interest from both service and content providers, the topic is expected to remain an open field for research and innovation for many years to come.

Acknowledgements

Maria Torres Vega is funded by the Research Foundation - Flanders.

References

[1] A. Bentaleb, B. Taani, A. C. Begen, C. Timmerer and R. Zimmermann, "A Survey on Bitrate Adaptation Schemes for Streaming Media Over HTTP," *IEEE Communications Surveys Tutorials,* vol. 21, no. 1, 2019.

[2] M. Seufert, S. Egger, M. Slanina, T. Zinner, T. Hossfeld and P. Tran-Gia, "A Survey on Quality of Experience of HTTP Adaptive Streaming," *IEEE Communications Surveys Tutorials,* vol. 17, no. 1, 2015.

[3] J. van der Hooft, M. Torres Vega, S. Petrangeli, T. Wauters and F. De Turck, "Tile-Based Adaptive Streaming for Virtual Reality Video," *Submitted to ACM Transactions on Multimedia Computing, Communications, and Applications,* 2019.

[4] M. Levoy and P. Hanrahan, "Light Field Rendering," in *Annual Conference on Computer Graphics and Interactive Techniques*, 1996.

[5] A. Clemm, M. Torres Vega, H. K. Ravuri, T. Wauters and F. De Turck, "Towards Truly Immersive Holographic-Type Communication: Challenges and Solutions," *Submitted to IEEE Communications Magazine*, 2019.

[6] M. Wijnants, H. Lievens, N. Michiels, J. Put, P. Quax and W. Lamotte, "Standards-Compliant HTTP Adaptive Streaming of Static Light Fields," in *ACM Symposium on Virtual Reality Software and Technology*, 2018.

[7] J. R. Daniel, B. Hernàndez, C. E. Thomas, S. L. Kelley, P. G. Jones and C. Chinnock, "Initial Work on Development of an Open Streaming Media Standard for Field of Light Displays (SMFoLD)," *Electronic Imaging,* vol. 2018, no. 4, 2018.

[8] P. A. Kara, A. Cserkaszky, M. G. Artini, A. Barsi, L. Bokor and T. Balogh, "Evaluation of the Concept of Dynamic Adaptive Streaming of Light Field Video," *IEEE Transactions on Broadcasting,* vol. 64, no. 2, 2018.

[9] Computer Graphics Laboratory, Stanford University, "The (New) Stanford Light Field Archive," 2008.

[10] S. Schwarz, M. Preda, V. Baroncini, M. Budagavi, P. Cesar, A. Chou, R. A. Cohen, M. Krivokuca, S. Lasserre, Z. Li, J. Llach, K. Mammou, R. Mekuria, O. Nakagami, E. Siahaan, A. Tabatabai, A. M. Tourapis and V. Zakharchenko, "Emerging MPEG Standards for Point Cloud Compression," *Journal on Emerging and Selected Topics in Circuits and Systems,* vol. 9, no. 1, 2018.

[11] E. d'Eon, T. Myers, B. Harrison and P. A. Chou, "ISO/IEC JTC1/SC29 Joint WG11/WG1 (MPEG/JPEG) input document WG1M40059/WG1M74006. 8i Voxelized Full Bodies - A Voxelized Point Cloud Dataset," 2017.

[12] M. Hosseini and C. Timmerer, "Dynamic Adaptive Point Cloud Streaming," in *Packet Video Workshop*, 2018.

[13] L. He, W. Zhu, K. Zhang and Y. Xu, "View-Dependent Streaming of Dynamic Point Cloud over Hybrid Networks," in *Advances in Multimedia Information Processing*, 2018.

[14] J. van der Hooft, T. Wauters, F. De Turck, C. Timmerer and H. Hellwagner, "Towards 6DoF HTTP Adaptive Streaming Through Point Cloud Compression," in *ACM Multimedia Conference*, 2019.

[15] S. Dijkstra-Soudarissanane, K. El Assal, S. Gunkel, F. ter Haar, R. Hindriks, J. W. Kleinrouweler and O. Niamut, "Multi-Sensor Capture and Network Processing for Virtual Reality Conferencing," in *ACM Multimedia Systems Conference*, 2019.

[16] F. Qian, B. Han, J. Pair and V. Gopalakrishnan, "Toward Practical Volumetric Video Streaming on Commodity Smartphones," in *International Workshop on Mobile Computing Systems and Applications*, 2019.

[17] M. Wien, J. M. Boyce, T. Stockhammer and W. Peng, "Standardization Status of Immersive Video Coding," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems,* vol. 9, no. 1, 2019.

[18] Google, "Draco," Google, 2016. [Online]. Available: https://github.com/google/draco. [Accessed 29 08 2019].

[19] R. Schnabel and R. Klein, "Octree-Based Point-Cloud Compression," in *Eurographics/IEEE VGTC Conference on Point-Based Graphics*, 2006.

[20] MPEG, "MPEG 3DG and Requirements - Call for Proposals for Point Cloud Compression V2," 2017.

[21] R. Verhack, T. Sikora, L. Lange, R. Jongebloed, G. Van Wallendael and P. Lambert, "Steered Mixture-of-Experts for Light Field Coding, Depth Estimation, and Processing," in *IEEE International Conference on Multimedia and Expo*, 2017.

[22] J. Du, Z. Zou, Y. Shi and D. Zhao, "Zero Latency: Real-Time Synchronization of BIM Data in Virtual Reality for Collaborative Decision-Making," *Automation in Construction,* vol. 85, 2018.

[23] M. Hosseini, "Adaptive Rate Allocation for View-Aware Point-Cloud Streaming," University of Illinois, 2017.

[24] J. Park, P. A. Chou and J. Hwang, "Rate-Utility Optimized Streaming of Volumetric Media for Augmented Reality," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems,* vol. 9, no. 1, pp. 149-162, 2019.

[25] Internet Engineering Task Force, "Identifying Our Deliverables," 2018. [Online]. Available: https://mailarchive.ietf.org/arch/msg/quic/RLRs4nB1lwFCZ_7k0iuz0ZBa35s. [Accessed 23 08 2019].

[26] Y. Cui, Y. Li, C. Liu, X. Wang and M. Kühlewind, "Innovating Transport with QUIC: Design Approaches and Research

Challenges," *IEEE Internet Computing,* vol. 21, no. 2, 2017.

[27] S. Petrangeli, D. Pauwels, J. van der Hooft, M. Ziak, J. Slowack, T. Wauters and F. De Turck, "A Scalable WebRTC-Based Framework for Remote Video Collaboration Applications," *Multimedia Tools and Applications,* vol. 78, no. 6, 2019.

[28] D. You, T. V. Doan, R. Torre, M. Mehrabi, A. Kropp, V. Nguyen, H. Salah, G. T. Nguyen and F. H. P. Fitzek, "Fog Computing as an Enabler for Immersive Media: Service Scenarios and Research Opportunities," *IEEE Access,* vol. 7, 2019.

[29] E. Bastug, M. Bennis, M. Medard and M. Debbah, "Toward Interconnected Virtual Reality: Opportunities, Challenges, and Enablers," *IEEE Communications Magazine,* vol. 55, no. 6, 2017.

[30] D. Szabo, A. Gulyas, F. H. P. Fitzek and D. E. Lucani, "Towards the Tactile Internet: Decreasing Communication Latency with Network Coding and Software Defined Networking," in *European Wireless Conference*, 2015.

[31] A. Javaheri, C. Brites, F. Pereira and J. Ascenso, "Subjective and Objective Quality Evaluation of 3D Point Cloud Denoising Algorithms," in *IEEE International Conference on Multimedia & Expo Workshops*, 2017.

[32] I. Viola, M. Rerábek, T. Bruylants, P. Schelkens, F. Pereira and T. Ebrahimi, "Objective and Subjective Evalation of Light Field Image Compression Algorithms," in *Picture Coding Symposium*, 2016.

Jeroen van der Hooft (S'14, M'18) obtained his M.Sc. and Ph.D. degrees in Computer Science Engineering from Ghent University, Belgium, in 2014 and 2019, respectively. He is currently active as a postdoctoral fellow at the Department of Information Technology, Ghent University – imec. His research interests include the end-to-end QoE optimization in adaptive video streaming and low-latency delivery of immersive video content.

Maria Torres Vega (S'14, M'17) obtained her M.Sc. degree in Telecommunication Engineering from the Polytechnic University of Madrid, Spain, in 2009, and her Ph.D. degree from the Eindhoven University of Technology, The Netherlands, in 2017. She is currently active as a postdoctoral fellow at Ghent University - imec. Her research interests include QoS and QoE in immersive multimedia systems and autonomous network management.

Tim Wauters (M'07) obtained his M.Sc. and Ph.D. degrees in Electrotechnical Engineering from Ghent University, Belgium, in 2001 and 2007, respectively. He is currently active as a postdoctoral fellow at Ghent University - imec. His work has been published in over 100 scientific publications in international journals and in the proceedings of international conferences. His research interests include network and service architectures, and management solutions for multimedia delivery.

Hemanth Kumar Ravuri obtained his B.Sc. degree in Electronics and Communication Engineering from the Jawaharlal Nehru Technological University, India, in 2014 and M.Sc. degree in Electrical Engineering with emphasis on Telecommunication Systems from BTH, Karlskrona, Sweden, in 2016. He is currently pursuing his Ph.D. degree at Ghent University - imec. His research interests include network and service architectures, next generation multimedia delivery, QoS and QoE.

Christian Timmerer (M'08, SM'16) obtained his M.Sc. and Ph.D. degrees from the Alpen-Adria-Universität Klagenfurt, Austria, in 2003 and 2006, respectively. He is currently active as an associate professor and vice-chair at the Institute of Information Technology at the same university. His research interests include immersive multimedia communication, streaming, adaptation, QoE and sensory experience. In 2013 he co-founded Bitmovin, where he is active as Chief Innovation Officer.

Hermann Hellwagner (S'85, A'88, M'95, SM'11) obtained his M.Sc. and Ph.D. degrees from the University of Linz, Austria, in 1983 and 1988, respectively. He is currently active as full professor at the Institute of Information Technology at the Alpen-Adria-Universität Klagenfurt. His research interests include multimedia communication and content adaptation, information-centric networking, and performance analysis of computer and communication systems.

Filip De Turck (S'95, M'98, SM'12) obtained his M.Sc. and Ph.D. degrees in Electronic Engineering from Ghent University, Belgium, in 1997 and 2002, respectively. He is currently active as a full professor at Ghent University - imec, where he leads the network and service management research group at the Department of Information Technology. His research interests include telecommunication network and service management, and design of efficient virtualized network systems.