SC22 Network Research Exhibition: Demonstration Preliminary Abstract

Uncompressed 8K video processing using edge-computing

Katsuhiro Sebayashi, Kanagawa Institute of Technology,
sebayashi@nw.kanagawa-it.ac.jp

Abstract

We aim to establish a video processing platform for uncompressed 8K ultra-high-definition videos, which can freely link transmission, storage, and processing functions and automatically configure the required video production workflow by using computing resources distributed among data centers and edges in the cloud.

In this NRE, we demonstrate an experiment that displays edited and processed videos at the venue by remotely using the edge devices of the video processing platform installed in Japan, which has a video processing capacity of 400 Gbps. This video traffic is visualized by two different real-time monitoring systems.

Overview

8K ultra-high-definition videos (8K videos) require 16 times more pixels than available in HDTVs, and are expected to be utilized in various industrial fields, such as video production and medicine. To enable 8K videos, it is essential to process and edit the video materials. However, uncompressed 8K videos are difficult to handle in real-time due to limitations in network bandwidth and device processing capacity. Therefore, we currently use a nonlinear editing machine that processes a video stored in local storage in non-real-time.

Our goal is to realize a video processing service that utilizes cloud and edge computing to enable real-time editing and distribution of 8K videos whenever needed, without requiring the users to have special equipment.

Cloud-based ultra-high definition (UHD) video processing uses a large amount of data that requires management and resource contention, making it difficult to guarantee high performance. Therefore, in this study, we propose a video processing platform architecture that can overcome these limitations and process large amounts of uncompressed 8K UHD video data with low latency by introducing an edge between the terminal and the cloud. Our proposed video processing platform realizes a series of video processing workflows by linking VVFs (Virtualized Video-handling Functions), as shown in Figure 1.

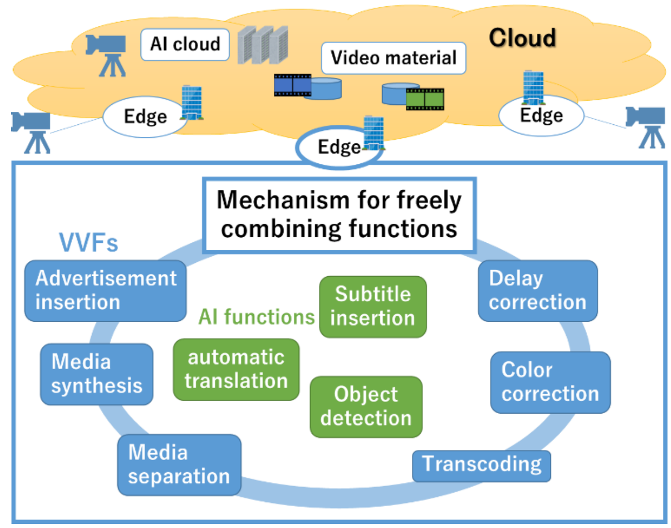

Figure 1. Cloud-based UHD video processing platform

Proposal of VVF

Each VVF applies the DPDK (Data Plane Development Kit) framework used in software router switches for video processing. The delay adjustment and video switching VVFs that process video data at the packet level can process four real-time streams (96 Gbps) of uncompressed 8K video data at 24 Gbps, as shown in Figure 2.

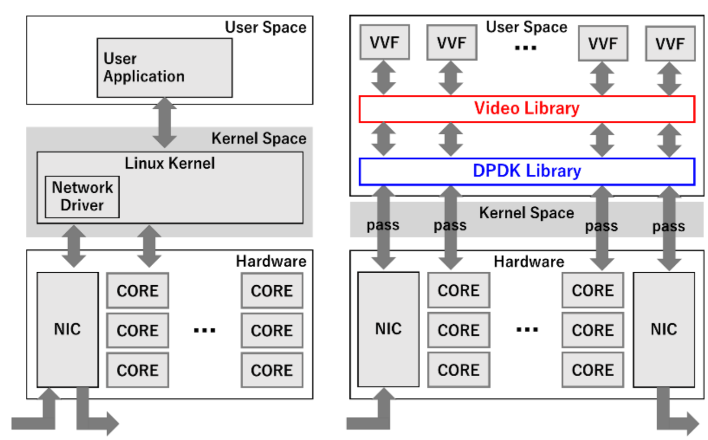

Figure 2. High-speed VVF using DPDK (right) and Normal DPDK usage (left)

Currently, 8K Dual Green (8K-DG) and full-resolution 8K (8K-YUV) devices are available on the market. Moreover, color adjustment is necessary to match the color tones when switching and stitching multiple video sources together for editing. Therefore, we developed a transcoding VVF, which converts from 8K-DG to 8K-YUV, including color correction. This VVF is not processed in packet (9000 byte) units, but in single frame (50 Mbyte) units, so memory bandwidth affects the processing time. Moreover, the YUV conversion process of transcoding is computationally demanding because it involves calculations, such as multiplication of a large amount of data.

The transcoding VVF implemented in conventional software took 200 ms to 240 ms per frame, and real-time processing could not be realized. Therefore, real-time processing was achieved by using SIMD extension instructions such as SSE (Streaming SIMD Extensions) and AVX (Advanced Vector Extensions) together with the DPDK framework, and by increasing the number of memory slots (number of channels) used, as shown in Figure 3.

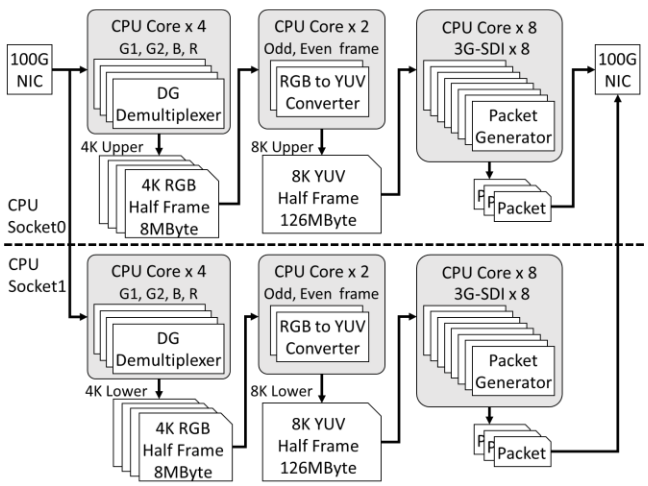

Figure 3. Implementation of transcode function

Prototype of the video processing platform and testbed

The video processing platform realizes a series of video processing workflows by linking VVFs that provide primitive video processing functions. We have built a prototype edge device comprising three PC servers and a 400G switch in the Sagamihara Data Center (DC) of SINET6, as shown in Figure 4.

Figure 4. Edge devices at Sagamihara DC

Furthermore, a testbed using the Ginoza DC of SINET6 in Okinawa, an 8K camera from Kanagawa Institute of Technology, and the 8K video servers of JGN's SterBED as video sources has been constructed to develop and verify the new VVFs, as shown in Figure 5.

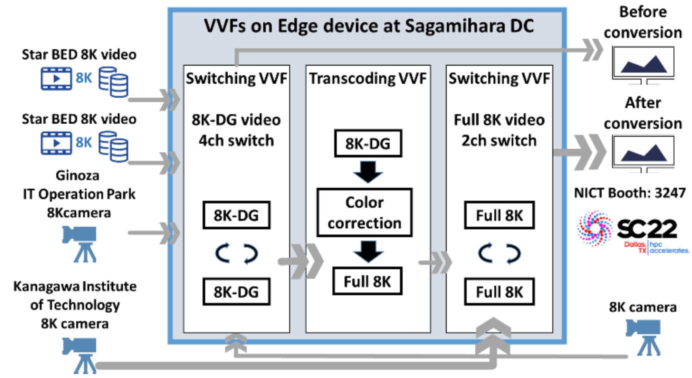

Figure 5. Uncompressed 8K UHD video workflow by chaining multiple VVFs

Additionally, to enable the provisioning of VVFs to servers in different DCs in the cloud, the application of SRv6-based service chaining for VVF linkage was also considered.

In February 2021, we constructed an SRv6 network between Tokyo and Osaka in JGN and successfully conducted a video-switching experiment, wherein the VVFs were linked by service chaining. This experiment verified the data plane, and studies are currently underway to include the control plane.

Monitoring

Traffic monitoring is important for the stable transmission of large-capacity real-time video streams such as uncompressed 8K videos. We are also developing two real-time monitoring systems with over 100 Gbps for multiple uncompressed 8K video streams, as shown in Figure 6.

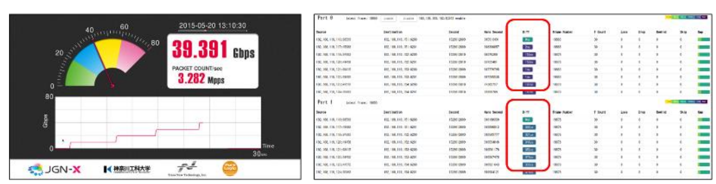

Figure 6. Screenshot of Traffic Mater (left) and Stream Monitor (right)

The first is the Traffic Meter, which visualizes the amount of traffic in real-time with a resolution of less than one second. It is possible to capture momentary video disturbances as changes in packet spacing (bursts).

The second is the Stream Monitor. A 24 Gbps 8K-DG video stream comprises 8 sub-streams of 3 Gbps, and a 48 Gbps 8K-YUV video stream comprises 16 sub-streams of 3 Gbps. For the video to be played back with no disturbance, these sub-streams must arrive with a delay difference within the acceptable range of the receiver's buffer capacity, with a frequency of order reversal and a missing rate. Our stream monitor lists the delay difference, reversal frequency, and missing rate for each sub-stream at one-second intervals, allowing the user to identify problematic sub-streams.

Demonstration

In this demonstration, a recently established edge device in Japan with a 400 Gbps video processing capacity will be operated remotely from the venue. Multiple 8K uncompressed video materials on JGN and SINET will be edited and processed by this edge device, with both the unedited and edited videos simultaneously displayed at the venue for demonstration purposes. This operation is performed via buttons and knobs on the MIDI Pad.

The 70-inch 8K display monitor displays color-adjusted and transcoded 8K-YUV video, which can be compared with the original video on the 28-inch monitor. The red, blue, and green colors can be adjusted with the knob, as shown in Figure 7.

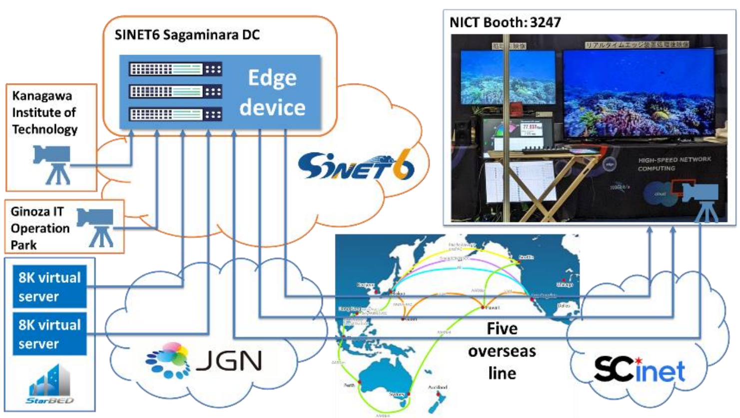

Figure 7. Network topology and traffic flow of the demonstration

The video traffic will be monitored in real-time by the Traffic Meter and Stream Monitor. With these monitoring screens, if there is a disturbance in the received video, the traffic behavior can be checked with the Traffic Meter, and the status of the sub-streams of the video layer can be checked with the Stream Monitor.

Goals

1. Simultaneous transmission of two types of 8K video: 8K-YUV (51.2 Gbps) after processing by the edge device and 8K-DG (25.6 Gbps) before processing, to visualize the video processing speed
2. Remote control of edge computing nodes in Japan from the United States
3. Monitoring the processing resources using a high-precision network monitoring device

Resources

- Stable connection of MTU9100 between Japan and the US through 80 Gbps bandwidth and IPV4 multicast traffic
- Tapping points or monitoring for debugging and troubleshooting

Involved Parties

- Mitsuru Maruyama, Kanagawa Institute of Technology, maruyama@nw.kanagawa-it.ac.jp
- Hajime Iwata, Kanagawa Institute of Technology, hajime@nw.kanagawa-it.ac.jp
- Hiroyuki Kimiyama, Daido University, kimiyama@daido-it.ac.jp
- Takayuki Nakachi, University of the Ryukyus, tnakachi@lab.u-ryukyu.ac.jp
- Kouta Aoki, Miharu Communications Inc., kota.aoki@miharu.co.jp
- Yasuhisa Kato, Miharu Communications Inc., y-kato@miharu.co.jp
- Shigeo Urushidani, National Institute of Informatics, urushi@nii.ac.jp
- Takeshi Kurimoto, National Institute of Informatics, tkurimoto@nii.ac.jp
- Eiji Kawai, National Institute of Information and Communications Technology, eiji-ka@nict.go.jp
- Hideki Otsuki, National Institute of Information and Communications Technology, eiji@nict.go.jp
- Kazumasa Kobayashi, National Institute of Information and Communications Technology, kazu-k@fatware.jp

Acknowledgements

The research results were partially obtained from NICT's commissioned research No. 03101.