SHD Digital Cinema Distribution over a Global High-speed Network: Internet2

Takahiro Yamaguchi†, Daisuke Shirai, Tatsuya Fujii, and Tetsuro Fujii

Abstract

We have developed a prototype SHD (Super High Definition) digital cinema distribution system that can store, transmit, and display motion pictures with eight million pixels per frame. In a long-distance SHD digital cinema transmission experiment performed on the Internet2 network in October 2002, we used an enlarged TCP (transmission control protocol) window, multiple TCP connections, and a shaping function to control the data transmission quantity. As a result, we succeeded in transmitting contents at about 300 Mbit/s between Chicago and Los Angeles, a distance of more than 3000 km. In this paper, we review the lessons we learned from this experiment.

1. Introduction

While cinema dominates the visual entertainment field, digital cinema projection using the High Definition TV (HDTV) is making major inroads. To date, however, the lower resolution of HD digital cinema has prevented it from matching the image quality of film [1]. The spatial resolution of 35-mm movie (original negative[*1]) film exceeds one thousand scanning lines [2]. Therefore, higher-quality digital cinema systems are desired [3]. One approach to addressing this problem is Super High Definition (SHD) images (resolution ≥2000 scanning lines), which surpass the quality of 35-mm film in terms of spatial resolution and approach the quality of 60-mm film [4].

Against this background, we have developed an SHD digital cinema system [5]-[7] that uses a 3840 × 2048 pixel image format with 30-bit color separation and 24 frames-per-second (fps) motion. The image resolution obtained with this system is about four times that of the 1080/24p format (1920 × 1080 pixels) used for HDTV, as shown in Fig. 1. Film-based movie contents are digitized by scanning 35- or 65-mm films, after which the resulting digital data is compressed and stored. Transmitting movie contents using public optical networks requires an exception-

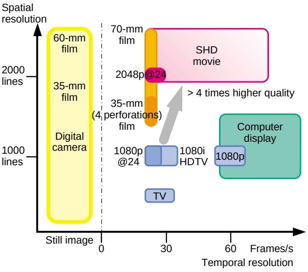

Fig. 1. Resolution and frame rate of SHD image.

† NTT Network Innovation Laboratories
Yokosuka-shi, 239-0847 Japan
E-mail: yamaguchi.takahiro@lab.ntt.co.jp

*1 Original negative: The original negative is also called the camera negative. There are three different basic types of 35-mm movie film: camera negative, intermediate, and print. Camera negative records as much detail as possible from the original scene, both spatially and in light range to make that original detail eventually available on a multitude of intermediate negatives from which are produced thousands of release prints for projection.

ally high-performance decoder and imaging system if the movies are to be displayed in real time as they are streamed. This is because the total bit rate of an SHD movie can be as high as 5.7 Gbit/s, so the movies need to be compressed by 10:1 or more to transmit them via envisioned wide-area optical IP networks using gigabit Ethernet (GbE). Additionally, an SHD real-time decoder and SHD projection device will be required at the SHD digital cinema. The high-speed optical networks will distribute the contents of VOD (video-on-demand) servers to wherever the SHD real-time decoders and projectors have been installed.

2. SHD digital cinema distribution system

Our SHD digital cinema distribution system, shown in Fig. 2, comprises three main devices: a real-time decoder, an SHD projector, and a movie server.

2.1 JPEG2000 real-time decoder

We have developed a hardware JPEG2000 decoder using an IA-32-based Linux personal computer (PC). Four JPEG2000 decoder boards are installed on the 64-bit PCI-bus of the PC. Each board has 30 JPEG2000 processors (Analog Devices Inc. ADV-JP2000) that process a quarter (1920 × 1024 pixels) of the whole image area. The decoder can decompress data in real time at a speed of 48 fps for 3840 × 2048 pixel 30-bit color images (i.e., 400 megapixels per second), using 120 JPEG2000 processors working in parallel. A control program running as an application consists of two threads with dual CPUs (two 1.44-GHz Pentium IIIs) that share the PC's main memory as a large data buffer. The decoder receives the JPEG2000-coded data stream from the movie server via the GbE network interface and decodes it quickly enough for real-time projection.

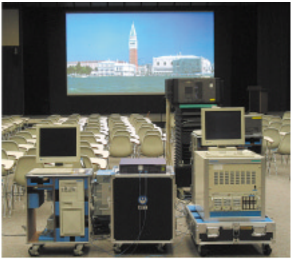

Fig. 2. Overview of SHD digital cinema distribution system. From the front left: server, GbE-switch, and real-time decoder; at the rear: projector.

2.2 SHD D-ILA projector

The SHD projector uses three prototype eight-megapixel (3840 × 2048 pixel) D-ILA devices (reflective LCD panels) produced by JVC, one for each 10-bit RGB (red, green, blue) color channel. Its 1600-W xenon lamp gives an effective brightness exceeding 3000 ANSI lumens, which is bright enough to show images on a 300-inch (diagonal) screen. We consider that the horizontal resolution of 3840 pixels will cover almost all image applications and enable us to fully utilize the vertical resolution of 2048 pixels even if the movie has a wide aspect ratio. The 30-bit color depth eliminates pseudo-edges that may be perceived on the gradation pattern of computer graphics based images. The high refresh rate (96 Hz) of the projectors completely eliminates flicker and is compatible with 24-fps movies. Every frame of the decoder output is simply displayed four times in the projector, without any interpolation between adjacent frames.

2.3 Movie server

The movie server is an IA-32-based Linux PC with dual CPUs (two 1.44-GHz Pentium IIIs) and has a high-performance RAID0*1 system. The movie content is digitized (3840 × 2048 pixels and 8- or 10-bit RGB per pixel) from 35-mm film by using the highest quality pin registered film scanner (IMAGICA Corp., ImagerXE). The film-digitized or computer-generated movie data is (1) divided into 128 × 128 pixel image tiles, (2) converted into YCbCr (4:2:2) color components, (3) compressed/encoded by a suitable data-format for the decoder, and (4) made into one big file bringing 1000 frames together as a unit that is stored in the server's RAID. The server program that receives the data transfer command reads the data from the RAID and then periodically writes it to the GbE network interface.

3. Internet2 digital cinema transmission trial

The fall 2002 Internet2 members meeting was held on October 28-29, 2002 at the University of Southern

*1 RAID: redundant arrays of inexpensive disks.
RAID0: striping mode.

California (USC). During this meeting, we performed an SHD digital cinema distribution (TCP/IP-based SHD digital cinema streaming) experiment linking Chicago and Los Angeles via the Internet2 network [8]. This was the world's first long-distance transmission of an SHD digital cinema data stream. We used the TCP/IP protocol to stream the digital cinema contents, but with this protocol the network throughput temporarily falls if one or more packets end up being dropped so that they must be re-transmitted. To maintain continuous projection if such a network stall occurs, the decoder spools the received data stream (4–8 s of movie data) in memory.

3.1 Network configuration

The network configuration is shown in Fig. 3. We set up the server in the Electronic Visualization Laboratory (EVL) at the University of Illinois, Chicago (UIC). The decoder and projector were installed at the Robert Zemeckis Center of the School of Cinema-Television at the University of Southern California (USC), Los Angeles. UIC and USC were connected to the Internet2 network. A bit stream of SHD movie-data, originating from the server at the EVL, was transmitted over the backbone network containing a GbE switch (OptiSwitch 4000 of MRV Communications, formerly Zuma Networks) etc., and projected at the Robert Zemeckis Center. The distance between the server and decoder was more than 3000 km and there were seven router hops between them. The round trip time (RTT) of the network was measured to be 59 ms.

3.2 Technique for long-distance distribution

We used (1) a large TCP window, (2) multiple TCP connections between the server and decoder, and (3) a shaping control function that depended on the data transmission quantity to improve the throughput of the long-distance TCP transmission. These are explained in detail below.

(1) Enlarged TCP window

The TCP window size is the amount of data that can be sent without acknowledgment. There are theoretical limitations to TCP-window-based flow control. The theoretical TCP throughput is

$$\text{TCP throughput} < \text{window size} \div \text{RTT}.$$

Therefore, the configuration guideline is

$$\text{window size} > \text{required TCP throughput} \times \text{RTT}.$$

Initially, the TCP window size of the system was extended to 4 MB from its initial 64 KB value. Accordingly, the theoretical throughput became about 540 Mbit/s for an RTT of 59 ms and window size of 4 MB. (For comparison, the theoretical throughput for a 64-KB TCP window is about 8.5 Mbit/s.) In a preliminary test, however, a stream transmission of 100 Mbit/s or more proved to be

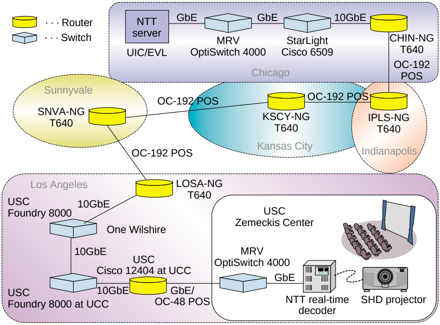

Fig. 3. Internet2 experimental network configuration.

impossible, although a stream transmission of 50 Mbit/s was possible.

(2) Use of multiple TCP connections

We transmitted streams while increasing the number of TCP connections between the server and decoder when the initial TCP window size was 4 MB. The server application divided the movie data into equal segments and sequentially wrote them to multiple TCP sockets connected to the decoder. Throughput increased with the number of TCP connections. Table 1 shows the success or failure for various bit rates and TCP connection numbers. A 200-Mbit/s stream transmission was possible with 64 TCP connections. However, it did not go up as expected from the theoretical value, and hit a ceiling at about 200 Mbit/s. The stream's bit rate could not be raised even when the number of connections was increased beyond 64. As the 200-Mbit/s movie data did not have the quality of the original/master 35-mm film, we had to find another way to improve the throughput further.

(3) Traffic shaping/smoothing

We used an application traffic monitor (a product of NTT Labs., available from Anritsu Co as MD1230A-20) to observe the traffic pattern with a resolution of 1 ms. The data transmission was very bursty, and the biggest peak exceeded 800 Mbit/s, as shown in Fig. 4. We expected that a momentary peak rate of 800 Mbit/s or higher would be a problem. A preliminary experiment that used UDP/IP (UDP: user datagram protocol) showed that packet loss would occur if the call transfer rate exceeded 800 Mbit/s in the Internet2 network.

To suppress the burstiness, we built a shaping control function for the data transmission into the socket writing process of the server application. The application divides the data stream after making a multiple TCP connection between the server and decoder, and writes the divided data segments into each socket gradually by adding a waiting time, as shown in Fig. 5. It was necessary to add a suitable waiting time when a single movie data frame was sent during a 1/24-s cycle. As a result, the large peaks could be eliminated, as shown in Fig. 6. We thus chose a TCP window size of 4 MB as the initial value, and transmitted movie data at an average

Table 1. Results of multiple TCP connection trials.

Number of TCP connections	Bit rate (Mbit/s): 50	100	150	200	250	300
1	○	×	—	—	—	—
4	○	○	×	—	—	—
16	—	○	○	×	—	—
64	—	—	○	○	×	×
80	—	—	—	○	×	—

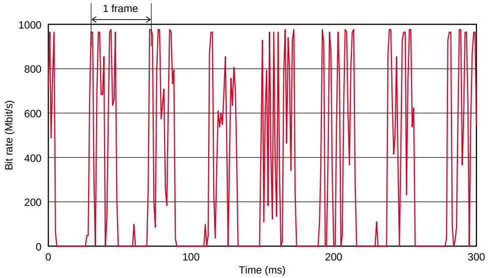

Fig. 4. Example of application traffic monitor outputs (without traffic shaping).

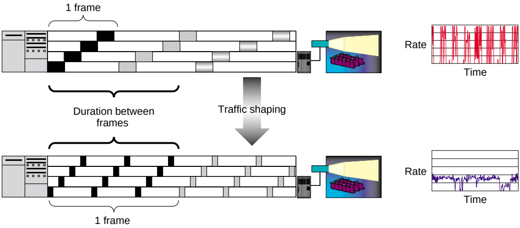

Fig. 5. Traffic shaping by server application.

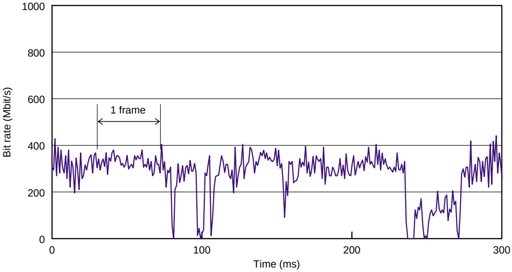

Fig. 6. Example of application traffic monitor output with traffic shaping.

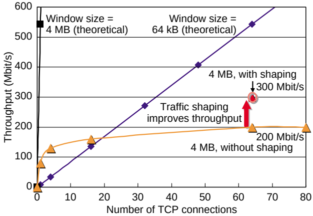

Fig. 7. Experimental results.

speed of 300 Mbit/s via 64 TCP streams with shaping control (indicated by the circle in Fig. 7).

3.3 Lessons learned

Figure 7 shows the results of the trials in regard to measures (2) and (3) described above. In using an SHD digital cinema for TCP transmission over high-speed, long-distance networks, we found that:

- a single TCP connection is likely to have poor performance
- the cause of the poor performance is packet loss (or potentially other causes)
- multiple TCP connections and traffic shaping let us transmit movie data successfully, but we are not completely sure why.

A complete analysis was impossible because of the lack of time (the Internet2 experiments lasted only a few days). However, the network's instability and poor controllability were probably the cause of the problems we encountered.

4. Conclusion

We have developed a digital cinema distribution system offering SHD image quality (3840 × 2048 pixels) that can transmit digital cinema contents with the quality of the original 35-mm film. The system requires high-quality connections between the server and decoder that can provide sustainable bandwidth without significant packet loss, delay, or jitter. The transmission of SHD movie data from a remote server to a distant decoder without interruption can only be accomplished over a long-distance network with sufficiently high capacity and quality of service. We performed a long-distance (i.e., long delay) streaming transmission experiment that involved three measures designed to improve transmission stability and successfully achieved stable SHD digital cinema streaming despite the long delay. The three measures were using a larger TCP window, multiple TCP connections, and a shaping function to control the data transmission quantity. The suppression of burstiness demonstrated in this trial won high praise from parties associated with the Internet2 Members Meeting and USC's School of Cinema-Television. The results we obtained boost the idea that SHD digital cinema can be distributed anywhere at any time.

Quantitative evaluations/analyses of the throughput in terms of RTT, TCP window size, and TCP connection number remain subjects of future study. We are currently refining the performance of the decoder to achieve higher image coding quality and to meet the various network bandwidth requirements.

Acknowledgment

We express our sincere gratitude to Professor Richard Weinberg of the University of Southern California for the support he provided us in conducting the experiments.

References

[1] http://www.screendigest.com/digitalcinema.htm
[2] R. R. A. Morton, M. A. Maure, and C. L. DuMont, "Assessing the Quality of Motion Picture Systems from Scene to Digital Data," SMPTE J, Vol. 111, No. 2, pp. 85-96, 2002.
[3] http://www.uemedia.com/CPC/article_5092.shtml
[4] S. Ono, N. Ohta, and T. Aoyama, "All-digital super high definition images," Signal Processing: Image Communication 4, pp. 429-444, 1992.
[5] T. Fujii, M. Nomura, J. Suzuki, I. Furukawa, and S. Ono, "Super High Definition Digital Movie System," SPIE VCIP'99, Vol. 3653, pp. 1412-1419, Jan. 1999.
[6] T. Fujii, M. Nomura, D. Shirai, T. Yamaguchi, T. Fujii, K. Hagimoto, and S. Ono, "IP Transmission System for Digital Cinema Using 2048 Scanning Line Resolution," IEEE GLOBECOM2002, pp. 1643-1647, Nov. 2002.
[7] T. Fujii, M. Nomura, D. Shirai, T. Yamaguchi, T. Fujii, and S. Ono, "Digital Cinema system using JPEG2000 movie of 8 million pixel resolution," IS&T/SPIE Electronic Imaging VICP2003, Santa Clara, California, U.S.A., 5022-07, Jan. 2003.
[8] http://oldwww.internet2.edu/apps/html/fall2002-demos.html#HDTVoverInternet2

Takahiro Yamaguchi
Research engineer, NTT Network Innovation Laboratories.
He received the B.E., M.E., and Ph.D. degrees in electronic engineering from the University of Electro-Communications, Chofu, Tokyo, in 1991, 1993, and 1998, respectively. He joined NTT Optical Network Systems Laboratories, Yokosuka, Japan in 1998 and has been working on research of super high definition imaging systems and their telemedical applications. He is a member of the Institute of Electronics, Information and Communication Engineers (IEICE) of Japan, the Institute of Image Information and Television Engineers (ITE) of Japan, and the Society for Information Display.

Daisuke Shirai
NTT Network Innovation Laboratories.
He received the B.E. degree in electrical engineering and the M.E. degree in computer science from Keio University, Hiyoshi, Kanagawa, in 1999 and 2001, respectively. He joined NTT Network Innovation Laboratories, Yokosuka, Japan in 2001 and has been researching a super high definition digital cinema system. He is a member of IEICE.

Tatsuya Fujii
Manager, Broadband Business Department, Business Communications Headquarters, NTT East Corp.
He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Tokyo, Tokyo, in 1986, 1988, and 1991, respectively. He joined NTT in 1991. He has been researching image processing and image communication networks. In 1996, he was a visiting researcher at Washington University in St. Louis. He is currently a member of IEEE, IEICE of Japan, and ITE.

Tetsuro Fujii
Group leader, media processing systems, NTT Network Innovation Laboratories.
He received the B.S., M.S., and Ph.D. degrees in electrical engineering from the University of Tokyo, Tokyo, in 1979, 1981, and 1984 respectively. He joined NTT in 1984. He has been researching adaptive acoustic signal processing, parallel digital signal processing, and image processing. He is a member of IEEE and IEICE.