High-Speed Immersive Media Stream Recorder

1. Research Team

Project Leader: Prof. Roger Zimmermann, *IMSC and Computer Science*

Other Faculty: Prof. Chris Kyriakakis, *Electrical Engineering*
Prof. Christos Papadopoulos, *Computer Science*

Graduate Students: Dwipal A. Desai, Kun Fu, Shihua Liu, Beomjoo Seo,

Industrial Partner(s): Intel

2. Statement of Project Goals

Presently, digital continuous media (CM) are well established as an integral part of many applications. In recent years, a considerable amount of research has focused on the efficient retrieval of such media [10, 13]. Scant attention has been paid to servers that can record such streams in real time. However, more and more devices produce direct digital output streams. Hence, the need arises to capture and store these streams with an efficient data stream recorder that can handle both recording and playback of many streams simultaneously and provide a central repository for all data. This project focuses on the design for a large-scale data stream recorder. Our goal is to introduce a unified architecture that integrates multi-stream recording and retrieval in a coherent manner. The project raises practical issues such as support for multizone disk drives, variable bit rate media, and disk drives that have a different write than read bandwidth.

3. Project Role in Support of IMSC Strategic Plan

The vision of the Integrated Media Systems Center is the development of immersipresence as the next great breakthrough in our digital era, moving people from the current two-dimensional world of computers, television and film to three-dimensional immersive environments with visual, aural and haptic capabilities [1]. One of its current instantiations is the Distributed Immersive Performance (DIP) effort. It is desirable to have recording, archiving and playback of performances. This requires a multichannel, multi-modal recording system that can store a distributed performance event in real-time, as it occurs. Ideally, such a system would allow the playback of the event with a delay that can range from very short to very long. We will pursue the design and implementation of such a recording system [2].

4. Discussion of Methodology Used

Figure 1 illustrates the architecture of a scalable data stream recorder operating in an IP network environment. Multiple, geographically distributed sources, for example video cameras, microphones, and other sensors, acquire data in real time, digitize it and send it to the stream recorder. We assume that the source devices include a network interface and that the data streams are transmitted in discrete packets. A suitable protocol for audio and video data traffic would be the Real-time Transport Protocol (RTP) while the client-recorder dialogue that includes control commands such as `start`, `record`, `pause`, `resume`, and `stop` is commonly handled via the Real-time Streaming Protocol (RTSP).

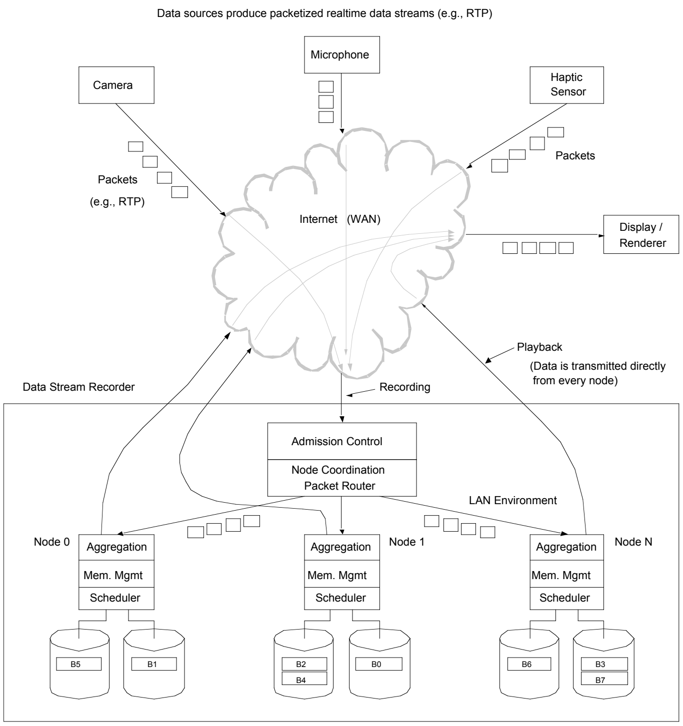

The data stream recorder receives the data with its *receiver* unit. This component performs the following functions: (1) admission control for new streams, (2) optional time stamping of packets, (3) packet-to-storage-node assignment and routing, and (4) storage node coordination and communication. The receiver forwards the incoming data to multiple *storage nodes*. Each storage node manages one or more local disk storage devices. The functions performed in the storage nodes are (5) packet-to-block aggregation, (6) memory buffer management, (7) block data placement on each storage device, (8) real time disk head scheduling, and (9) retrieval scheduling for outgoing streams.

A brief summary of the techniques employed are presented in the following list:

- Multi-node, multi-disk cluster architecture to provide scalability [2, 9].
- Random data placement for the following operations: block-to-storage node assignment, block placement within the surface of a single disk and optionally packet-to-block assignment. These result in the harnessing of the average transfer rate for multizone disk drives and improves scalability [12].
- Unified model for disk scheduling: deadline-driven data reading and writing (fixed block sizes reduce complexity of file system).
- Unified memory management with a shared buffer pool for both reading and writing.
- Selective retransmission based error control and bandwidth smoothing through feedback [6, 8, 11].
- Statistical admission control to accommodate variable bit rate (VBR) streams and multizone disk drives.

5. Short Description of Achievements in Previous Years

This project was started just recently (2002). Therefore, the accomplishments up to date are presented in Section 5a.

5a. Detail of Accomplishments During the Past Year

During the last year we have concentrated on the initial design of our stream media recorder. Some of the highlights of the past year were as follows:

- We have published a design paper [2] that presents the overall architecture of a large-scale data stream recorder. It discusses many of the inherent challenges that need to be solved.
- We have designed and analyzed a novel, high performance statistical stream admission control algorithm that considers a mix of reading and writing streams [14]. Our preliminary results indicate that our algorithm outperforms existing techniques by up to 50%.
- We have implemented a camera interface that acquires digital video from a FireWire (IEEE 1394) port in MiniDV format (640x480 pixels at 30 frames per second) and transmits the resulting data stream over an IP network with the RTP protocol. We will be using this interface to further test our stream recorder prototype.
- On Saturday, January 18, 2002, we used a preliminary prototype of our stream media recorder to record a multichannel audio performance over the Internet from New World Symphony' Lincoln Theater in Miami Beach. The cross-continental, low-latency streaming was executed with audio streaming software developed by Professor Papadopoulos' group. The eight channels of audio data were then forwarded to our recording system.

Up-to-date information can be found on our home page at http://dmrl.usc.edu/research.html.

6. Other Relevant Work Being Conducted and How this Project is Different

The goal of this project is the design of a unified systems architecture. The Multicast Multimedia Conference Recorder (MMCR) [3] probably is the most related to our work. The purpose of this project was to capture and play back multicast (MBone) sessions. The authors list a number of interesting and relevant issues for such systems. They focus more on the higher level aspects such as indexing and browsing the available sessions, while assuming only a small number of concurrent sessions. Our design, on the other hand, is specifically concerned with a scalable, high performance architecture where resources (memory, disk space and bandwidth) need to be carefully scheduled. There are also commercial systems available that relate to our design. We classify them into the following three categories.

1. Streaming media systems (e.g., Microsoft's Windows Media, Apple's QuickTime, RealNetworks' RealOne). These systems are optimized for streaming of previously (offline) stored content. Some of them also allow real time live streaming (i.e., forwarding with no recording). They are designed for multi-user access and multiple media types. They cannot usually take advantage of a cluster of server nodes.

2. Personal Video Recorders (PVR), for example the TiVo models and SonicBlue's ReplayTV. These systems allow real time recording and playback of standard broadcast quality video. Some of the limitations are that they are designed as single-user systems. Furthermore, they are optimized for a single media type (NTSC/PAL/SECAM video with two channels of audio). Local playback is supported and with newer models file sharing is enabled over a network. However, they do not provide streaming playback over a network.

3. Editing systems and broadcast servers. These systems are the professional cousins of the PVRs. They are used for the production and distribution of video content (e.g., to TV stations) and they are designed to interface via professional I/O standards (usually not Ethernet). Their use is for local environments, not distributed streaming setups. Most of the time they handle only a few media types and one (or a few) streams at a time. Their special purpose hardware and elaborate control interfaces to other studio equipment places them into a price category that makes them not cost-effective for use as a more general purpose stream recorder.

As indicated, none of these categories encompasses the full functionality that we envision. Each one of them only provides a subset of the desired functionalities.

7. Plan for the Next Year

This coming year we plan to extend the Remote Media Immersion (RMI) system to include a two-way (interactive) version. The most advanced scenario is in distributed (many-to-many) immersive activities, such as a distributed concert performance in which musicians located in physically different sites must coordinate their activities. This concept is termed a Distributed Immersive Performance (DIP). We plan to implement an initial stream recorder prototype based on our design concepts that will allow recording, archiving and playback of performances. In collaboration with the New World Symphony (Miami Beach, FL) we intend to perform a number of experiments to demonstrate our advanced recording technology.

8. Expected Milestones and Deliverables

Next year we plan to continue our work on the high-speed immersive media stream recorder such that it can support the requirements of the Distributed Immersive Performance project. Specifically, the following key issues will be addressed:

- **Initial prototype implementation.** It is desirable to have recording, archiving and playback of performances. This requires a multichannel, multi-modal recording system that can store a distributed performance event in real-time, as it occurs. Ideally, such a system would allow us to playback the event with a delay that can range from very short to very long. We will pursue the design and implementation of such a recording system.
- **Performance evaluation**. Together with several other IMSC investigators we will pursue a number of experiments to test and evaluate our first recorder implementation. Cross-continental streaming trials together with the New World Symphony will be among the tests.
- **Capabilities extension**. We will continue to investigate additional research issues that are not currently addressed in our initial design. Among them are the main memory buffer management with a unified pool of buffers and the write deadline determination of in-coming data. These ideas will require further studies to achieve our goal: that the overall performance of the system is maximized.

9. Member Company Benefits

The high-speed immersive media stream recorder prototype is being built partially with equipment donated by Intel.

10. References

[1] D. McLeod, U. Neumann, C.L. Nikias and A.A. Sawchuk. Integrated Media Systems. *IEEE Signal Processing Magazine*, vol. 16, no. 1, pp. 33-76, January 1999.

[2] R. Zimmermann, K. Fu, and W.-S. Ku, "Design of a Large Scale Data Stream Recorder," in Proceedings of the *5th International Conference on Enterprise Information Systems* (ICEIS 2003), (Angers, France), April 23-26 2003.

[3] L. Lambrinos, P. Kirstein, and V. Hardman. The Multicast Multimedia Conference Recorder. In *Proceedings of the International Conference on Computer Communications and Networks, IC3N*, Lafayette, Louisiana, 1998.

[4] R. Zimmermann and S. Ghandeharizadeh, "HERA: Heterogeneous Extension of RAID," in Proceedings of the *2000 International Conference on Parallel and Distributed Processing Techniques and Applications* (PDPTA 2000), (Las Vegas, Nevada), June 26-29 2000.

[5] R. Zimmermann, K. Fu, C. Shahabi, S.-Y. D. Yao, and H. Zhu, "Yima: Design and Evaluation of a Streaming Media System for Residential Broadband Services," in *VLDB 2001 Workshop on Databases in Telecommunications* (DBTel 2001), (Rome, Italy), September 2001.

[6] R. Zimmermann, K. Fu, N. Nahata, and C. Shahabi, "Retransmission-Based Error Control in a Many-to-Many Client-Server Environment," in Proceedings of the *SPIE Conference on Multimedia Computing and Networking 2003* (MMCN 2003), (Santa Clara, California), pp. 34–44, January 29-31 2003.

[7] R. Zimmermann, "Streaming of DivX AVI Movies," in Proceedings of the *ACM Symposium on Applied Computing* (SAC 2003), (Melbourne, Florida), March 9-12 2003.

[8] R. Zimmermann, K. Fu, M. Jahangiri, and C. Shahabi, "A Multi-Threshold Online Smoothing Technique for Variable Rate Multimedia Streams," Under review by the *Multimedia Tools and Applications* journal, 2003.

[9] C. Shahabi, R. Zimmermann, K. Fu, and S.-Y. D. Yao, "Yima: A Second Generation Continuous Media Server," *IEEE Computer*, vol. 35, pp. 56–64, June 2002.

[10] A. E. Dashti, S. H. Kim, C. Shahabi, and R. Zimmermann, eds., Streaming Media Server Design. Prentice Hall IMSC Press Multimedia Series, March 2003. ISBN: 0-130-67038-3.

[11] Ch. Papadopoulos and G. M. Parulkar. Retransmission-based Error Control for Continuous Media Applications. In Proceedings of the 6th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV 1996), Zushi, Japan, April 23-26 1996.

[12] Ashish Goel, Cyrus Shahabi, Shu-Yuen Didi Yao, and Roger Zimmermann. SCADDAR: An Efficient Randomized Technique to Reorganize Continuous Media Blocks. In Proceedings of the 18th *International Conference on Data Engineering* (ICDE 2002), San Jose, California, February 26-March 1, 2002.

[13] R. Zimmermann, M. R. Kolahdouzan, and C. Shahabi, "MASAA: A Case Study in Building a Distributed Integrated Media Database," in Proceedings of the *33rd Hawaii International Conference on System Sciences* (HICSS-33 '00), Maui, Hawaii, January 2000.

[14] R. Zimmermann and K. Fu, "Pushing the Limit of Resource Utilization: Statistical Admission Control for Streaming Media Servers," submitted for publication, 2003.