

Performance Evaluation of Media Losses in the Continuous Media Toolkit*

Duminda Wijesekera, Srivatsan Varadarajan, Shwetal Parikh,
Jaideep Srivastava and Anil Nerode[†]

Department of Computer Science, University of Minnesota, Minneapolis, MN 55455.
Institute for Foundations of Intelligent Systems, Cornell University, Ithaca, NY 14853[†].
e-mail: {wijesek|varadara|ssparikh|srivasta}@cs.umn.edu,
nerode@hybrid.cornell.edu

Abstract

This paper shows a methodology for user-driven, top down approach to research in quality of service issues in multimedia systems. As a case study, we show a development of metrics, validation by means of a user study, and a performance evaluation of a prototyping environments. What is used, the Berkeley Continuous Media Toolkit (CMT) is a popular environment that satisfies this need. From a human user's perspective, in order for multimedia demonstrations to be comprehensible, the number of audio or video frames dropped and the timing delays in the ones that are displayed, need to be kept to a minimum. Therefore, it is important to know the frame dropping characteristics of CMT. In a series of experiments we monitored the variation of these parameters with respect processor and network loads. It was observed that loads affect aggregate frame drops at lower rates and consecutive frame drops at higher rates. Because at a higher rates a large number of consecutive frames are dropped, the ones that are played appear in a more timely manner. As a solution to observed problems, we present some QoS based approaches to control drop and delay parameters.

Key Phrases: Quality of Service, User Centric QoS Research, Continuous media, Performance evaluation

1 Introduction

Traditionally human sensory organs process visual, aural, olfactory, taste and touch information. These correspond to the dimensions in which we comprehend the external world. The last six thousand years have used textual information – which until a decade ago was the only *medium* or *type* of information processed by the Von Neumann type of machinery. Thus naturally, the addition of our natural types of *information*

media or types of information to human interfaces of computing machines, lead to better representations and consequently quicker comprehension of the communicated message. Consequently, such media should be presented in a way that is in the wide range constituting from being very appealing to being marginally satisfactory manner for human consumers. This wide margin has been attempted to be quantitatively captured by a *Quality of Service (QoS)* metrics at the application level. Between mostly qualitatively stated application requirements and an acceptable high quality multimedia testbed, there are many stages of design, testing, performance evaluation and implementation. This paper describes a process that has been gainfully used in our research at the University of Minnesota and describes results of a prototyping effort on the *Berkeley Continuous Media Toolkit (CMT)* [SRY93, MPR97].

The rest of the paper is organized as follows: Section 2 presents our approach to QoS based multimedia testbed development, along with a summary of results of a case study in applying our methodology. The next couple of sections present various steps of this methodology applied to our case study: Namely, Section 4 presents our metrics for continuity, along with results of a user study to validate our metrics. Then, Section 5 presents the design and functionality of CMT, a testbed chosen to carry out the benchmarking part of our methodology. Section 7, 8 presents detailed results of our performance evaluations and their conclusions. Finally, Section 9 contains concluding remarks.

The general trend of our performance evaluations indicate that as the speed increases, the frame drops increase. Higher loads induce aggregate drops at lower speeds and consecutive drops at higher speeds. Furthermore, when a large number of frames are dropped by CMT, the remaining are displayed in a timely manner.

*This work is partially supported by Air Force contract number F30602-96-C-0130 to Honeywell Inc, via subcontract number B09030541/AF to the University of Minnesota and DOD MURI grant DAAH04-96-10341 to Cornell University

Definition	Source
The totality of features and characteristics of a product or service that bears on its ability to satisfy given needs.	[Ame78] [Deu] [Eur81]
Fitness for use and conformance to requirements	[Jur74]
Conformance to requirements	[Cro79]
The degree to which a product characteristics confirm to the requirements placed upon that product, including reliability maintainability and safety.	[Swi81]
The degree to which a product or service is fit for the specified need.	[Seg81]
User's subjective wishes or satisfaction with the quality of application performance -performance, synchronization cost, etc.	[VKvBG95]

Table 1: Definitions for Quality of Services/Products

2 A Methodology for User Based QoS Research

The fact that most multimedia is for human consumption, and due to inherent limitations of human perception, some loss of quality can be tolerated, can be utilized as a advantage for building multimedia systems. Consequently, given the realization that it is sufficient to provide multimedia services to be within such tolerable limits, the human consumer and multimedia systems can be conceived to be in a *client - service provider* relationship, where the qualitative expectations of the former need to be satisfied by the service provided by the latter. The generic term used to express such expectations quantitatively is *Quality of Service (QoS)*. One of the main tasks of this dissertation has been to provide some metrics for a class of multimedia services.

3 Our Approach to QoS Research

Before the advent of multimedia, consumers of a commodity or users of a service use to express their needs qualitatively as *quality of service requirements (QoS)*. Examples of such a QoS requirement will be the number of times a hinge of a door could be used before it wears out, or the strength of cable. Quality of services/products has been used in other disciplines of computing, such as networks [Tow93] and quality control [MVCC92] etc. According to [MVCC92], [SW85] QoS has been defined as given in Table 1.

Contents of table 1 indicates that most definitions for QoS are qualitative. Consequently, there is a need to make them quantitative. More importantly, for the systematic development healthy research program in QoS one also needs a methodology and an approach. The approach we developed for this process is diagrammatically given in Fig. 1. As the diagram indicates, the plan begins with identifying an aspect of the *real world* that will become the focus of study for QoS research. To measure *qualitative* requirements, metrics need to be defined for the chosen aspect. To test the relevance of

defined metrics, a user/application *validation* needs to be carried out, and until found relevant the metrics need to be revised. The next step is the evaluation of existing test-beds against chosen metrics. This step is necessary because, if some existing systems provide acceptable or higher quality services, then there is no need for developing new systems. The researcher in this case must enlarge the scope of the study, i.e. model a broader aspect of the real world, and repeat the previous steps. Conversely, if existing systems do not provide satisfactory services, then new systems need to be designed, evaluated and improved until they are found to provide satisfactory quality of services.




Figure 1: Our Approach to Quality of Service Research

Another issue in QoS research is that, most systems have a layered architectures, and thus any given layer is a client of the layer below it, and a service provider for those above it. Impact of this design on QoS research is two-fold: Firstly, QoS management belongs to several levels, and for any level to do QoS management, the levels below it must provide certain guarantees. This leads to interesting issues in QoS translation between layers. The second is that in discussing QoS the service consumer and the provider need to be clearly identified, just as in the ISO architecture for networks.

4 Continuity Metrics

For the purpose of describing QoS metrics for lossy media streams, we envision a CM stream as a flow of data units (referred to as logical data units - LDUs in the uniform framework of [SB96]). In our case, we take a video LDU to be a frame, and an audio LDU to constitute 8000/30, i.e. 266 samples of audio¹. Given a rate for streams consisting of these LDUs, we envision that there is time slot for each LDU to be played out. In the ideal case a LDU should appear at the beginning of its time slot.

4.1 Metrics for Stream Continuity

Continuity of a CM stream is measured by three components: *flow rate*, *timing drift* and *content loss*. The ideal rate of flow and the maximum permissible deviation from it constitute our *rate* parameters. As stated, given the ideal rate and the beginning time of a CM stream, there is an ideal time for a given LDU to be displayed. The appearance time of a given LDU may deviate from this ideal. Our *drift* parameters specify aggregate and consecutive non-zero drifts from these ideals, over a given number of consecutive LDUs in a stream. For eg., first four LDUs of two example streams, with their expected and actual times of appearance, are shown in Fig. 2. In the first example stream, the drifts are respectively 0.0, 0.2, 0.2 and 0.2 seconds; and accordingly it has an aggregate drift of 0.6 seconds per 4 time slots, and a non-zero consecutive drift of 0.6 seconds. In the second example stream, the largest consecutive non-zero drift is 0.3 seconds and the aggregate drift is 0.5 seconds per 4 time slots. The reason for a lower consecutive drift in stream 2 is that the unit drifts in it are more spread out than those in stream 1.

In addition to timing and rate, ideal contents of a CM stream are specified by the ideal contents of each LDU. Due to loss, delivery or resource over-load problems, appearance of LDUs may deviate from this ideal, and consequently lead to discontinuity. Our metrics of continuity are designed to measure the average and bursty deviation from the ideal specification. A loss or repetition of a LDU is considered a unit loss in a CM stream. (A more precise definition is given in [WS96].) The aggregate number of such unit losses is the *aggregate loss* of a CM stream, while the largest consecutive non-zero loss is its *consecutive loss*. In the example streams of Fig. 2, stream 1 has an aggregate loss of 2/4 and a consecutive loss of 2, while stream 2 has an aggregate loss of 2/4 and a consecutive loss of 1. Once again, the reason for the lower consecutive loss in stream 2 is that its losses are more spread-out than those of stream 1.

¹SunAudio has 8-bit samples at 8kHz, and an audio frame constitutes 266 such samples, equivalent to a play time of one video frame, i.e. 1/30 seconds.

4.2 Parameters for Continuity Metrics

In a user study [WSNF97] it has been determined that aggregate losses of upto 17/100 were imperceptible, those beyond 23/100 were annoying, and in between these two values were tolerable. The tolerable value for consecutive losses was determined to be two frames, i.e about 8000 samples. For audio this limit was about three frames. Due to the inability of precisely introducing timing drifts, the user study did not estimate any values for it.

5 CMT: Design and Functionality

Rapid growth of multimedia systems, and accordingly research efforts in this area, have made it necessary to have toolkit support for rapid prototyping. Being one of the most popular toolkits offered in multimedia, the *Berkeley Continuous Media Toolkit (CMT)* [SRY93, MPR97] has gone a long way in fulfilling this need. In [MPR97], Mayer-Patel and Rowe measured the performance of CMT and its system overheads. The Berkeley Continuous Media Toolkit, commonly referred to as CMT, has been constructed for the purposes of application development and research in multimedia systems [MPR97]. It is an extensible system consisting of a three layered architecture. The topmost layer, referred to as the *application code layer*, consists of Tcl, Tk and Tcl-Dp [SRS93] code. The next layer, consists of a CM object model, time and synchronization services, CM event services and storage/buffer services for CM streams. The bottom most *resource layer* consists of operating system and hardware services. The complete design of CMT is explained in detail in [MPR97].

One of the outstanding features of CMT is the ease of programming in Tcl/Tk, and the extensible and relatively *policy free* nature of implementation, which makes it a great testbed for experimentation.

CMT envisions multimedia as streams flowing in and out of CM objects, in their journey between *sources* and *sinks*. These CM objects correspond to *services* that are used by CM streams, such as *segments*, that represent data sources stored in files. Similarly, *play* objects represent stream players, such as speakers and video displays. CMT provides distributed multimedia services in the sense that objects in a CM pipeline can sit on different locations, thereby requiring network transport services. The network twin objects *packet source* (`pktSrc`) and *packet destination* (`pktDest`) provide these services of sending and receiving streamed data.

A CMT application program specifies pipelines of CM streams flowing between CM objects in Tcl/Tk. The objects can be created within some *CM process*. These scripts are interpreted by a Tcl interpreter extended with Tcl-DP to provide distributed services. For an

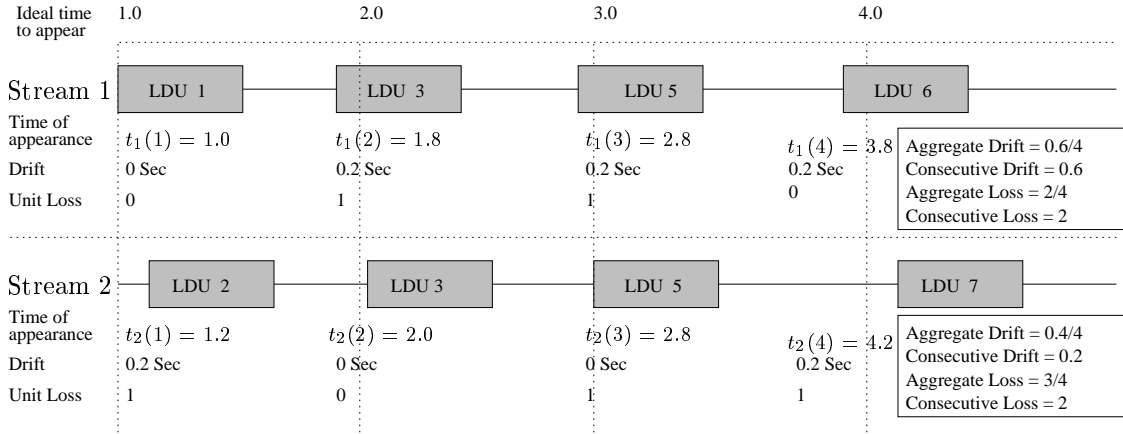


Figure 2: Two Example Streams used to Explain Metrics

application that uses only one location, a single CM process is created with appropriate sources and sinks. Distributed applications have CM processes at each location, where objects belonging to that location are created within the corresponding process.

The Logical Time System (LTS) in the CMT Library layer provides a mechanism for applications to maintain a concept of where in time the application is, and how fast time is progressing. More than one LTS can exist in one application and different CM processes can be synchronized by using the same LTS. CMT ensures that the LTS progresses according to the *rate* specified by the application program.

CMT maintains CM streams by passing data between objects by either the push model or the pull model. In our experimental setup, we used the push model, where the producer of data, typically the source object, initiates the data transfer to the object immediately downstream from it. The source object maintains a continuous streams of CM data by periodically invoking itself to fetch data from the specified file at the specified rate, and transfers data to the corresponding object downstream. The period of such data fetches are specifiable in application code. Objects that are being called by some other object upstream provides an `accept` method, that is being passed a buffer containing appropriate data by the callee. After *processing* the data in the buffer, such an intermediate object calls the `accept` method of its immediate downstream neighbor, thereby maintaining the flow of CM data through the specified pipeline, until a sink object such as a display device is reached. The sink objects are the final consumers of CM data, whereby it is displayed to the human viewer.

5.1 Application Code Used in Our Experiment

In our experiments, we use two application programs, referred to as *local playback application* and *remote play-*

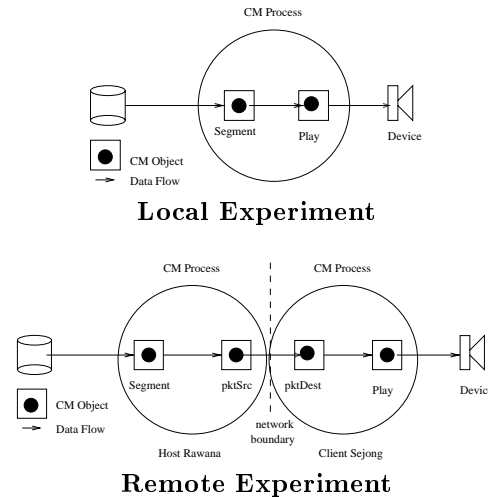


Figure 3: Local and Remote Playback Experiments in CMT

back application in [MPR97], graphically represented in Fig 3. The local experiment consists of a single CM process containing four objects, namely: `mjpegSegment`, `auSegment`, `mjpegPlay` and `auPlay`. As shown in Fig 3, they are pipelined so that `mjpegSegment` passes a video stream to `mjpegPlay` to be displayed at the local screen and `auSegment` passes an audio stream to `auPlay` to be displayed at the speaker. Transparent to the application programmer, the play objects call an `accept` method of a low level object called `device`, that is responsible for the final display of media units.

In the *remote play* experiment, the `segments` and `play` objects are on two different machines that are connected by Ethernet. Consequently, in addition to the objects used in the local play, *packet source* and *packet destination* network twins are used.

Because the performance analysis entails the behavior of objects used in our application code, following sections are devoted to a description of relevant properties of these objects.

5.1.1 Segment Object

Given the current value of the LTS object, the `segment` object determines the appropriate video frames or parts of audio frames to be fetched and played in the current fetch cycle. It then schedules itself to be called back after a lapse of the *cycle time*. The application code can also set a parameter *send ahead*, that specifies the time that the server is supposed to be ahead of the client.

In the process of being recalled, based on the time it was actually supposed to be called and the time it was called, the video `segment` determines the number of frames to be dropped by using an *inverse binary ordering* scheme [Smi]. This scheme attempts to minimize the consecutive dropping of frames, so that at the end, the stream would not appear *jittery*. In dropping the frames, the `mjpegSegment` also interpolates the display times of frames adjoining the dropped frame so that, although the frames were dropped, on the time scale no time gaps appear on the stream.

5.1.2 Play Object

In the application code we used, the play object checks and discards frames that are too late to be displayed, and then displays them in order. The video play object provides hooks to select the appropriate frame out of a collection that has been handed over to itself in a buffer.

5.1.3 Device Object

These objects are transparent to the programmer and contain low level calls specific to the platform and they submit data to be played, as has been called by the play object. The `SparcAuDevice` pre-empts *leftover* audio to *re-synchronize* at a regular frequency.

5.1.4 Network Twins

The network twins, `packet-Source` (`pktSrc`) and `packet-destination` (`pktDest`) uses a *cyclic UDP* protocol, where it is responsible for its own retransmission of the lost/delayed packets. This protocol has been explained in detail, with experimental verification of appropriateness for CM transport on best effort networks, in [Smi].

6 Experimental Evaluation of Media Losses

As stated, the objective of our experiments is to measure the performance of CMT with respect to continuity losses in audio and video. Our performance evaluation, is carried out with respect to metrics set forth in Section 4.

6.1 Types of Experiments

We carried out two main experiments. In the first one the segment object and the play object reside on the same machine, which we refer to as *local* objects. In the second experiment they are distributed across two machines. By the very design of CMT, this setup requires having *network twins*, the *packet source* and the *packet destination*. In the nomenclature of [MPR97], they are the *local playback application model* and the *simple remote playback application model*, for which the stream flows are given in Fig. 3.

In our experiments we measured the performance of CMT with respect to selected metrics under varying processor loads and network loads. Accordingly, for local experiments we collected data under different processor loads, and for remote experiments we collected data under different network and processor loads. In repeating experiments, the processor loads were measured by `top`, which gives the number of processes waiting for the processor, i.e. the length of the waiting queue. The experiments were carried out only when the length of the processor queue had stabilized for the immediately preceding 1, 5 and 15 minutes. The numbers reported are these stable values. The network loads were measured by the use of a *network sniffer*, and the experiments were carried out under stable loads reported by the sniffer in terms of packets per second.

6.2 Experimental Parameters and Methodology

In all of our experiments, we displayed audio and video streams, of length 1600 frames each. We ran the local experiment with different processor loads of 0, 1 and 5, which we call *low*, *medium* and *high* loads. For network loads we ran the experiment with approximately 0, 200 and 2000 packets/second. For each load setting we repeated our experiment five times and collected average statistics over them.

Our experiments were carried out using two uni-processor Ultra Sparc 1's with 64 MB RAM each, connected by a 10 Mbps Ethernet. Our work stations were equipped with JPEG Parallax[©] cards, and we used the Motion-JPEG video format and the 8-bit SparcAu format for audio.

As stated in [MPR97], some care must be taken to ensure that the experimental methodology does not degrade the performance of the system, i.e. it is as *non-intrusive* as possible. In this respect, we ensured that the time recorded by each module has been taken from already existing variables, and written into global arrays that were allocated before the system started to run in a stable state. Furthermore, these arrays were written into files when the objects were destroyed, thereby incurring minimal overhead in maintaining performance statistics. We also eliminated some initial readings to ensure that any start-up overhead does not affect fi-

nal measurements. Finally, in order to keep track of frame numbers, we have introduced frame numbers into header portions of video and audio frames. Statistics maintained were the frame number and the time of arrival at each relevant CM object.

7 The Local Experiment

This section presents results of the local experiment. i.e. the losses of audio and video streams where the `segment` and the `play` object reside on the same machine, and consequently there is no network involved. Our experiments were conducted with differing processor loads at speeds of 30 fps (speed = 1) to 300 fps (speed = 10) for both audio and video. The results are tabulated in terms of our metrics, i.e. aggregate and consecutive losses and drifts, for both media types.

Based on our experimental results we notice the following behavior of CMT.

1. The aggregate and consecutive frame losses increase with respect to speed for all video objects and audio segments.
2. Higher loads affect aggregate losses at low speeds and consecutive losses at higher speeds for all video objects.
3. Processor loads have no effect on audio at the `segment`. We believe that this is due to two conditions: (1) Audio data is relatively small. (2) There is no *delay based drops* for the audio stream at the `segment`, whereas in the video `segment`, drops are based on an *inverse binary order schema*.
4. At higher loads, `play` and `device` objects drop more frames than `segment` object. This is due to the fact, that they are invoked late, thereby drop frames.
5. As the speed increases, drifts drop. That means, as the speed increases, there are fewer frames present, but the ones that aren't dropped appear in time within their slot.
6. At higher loads and low speeds, audio `device` and `play` show larger drifts, indicating jitter.

8 Remote Experiment

This section presents results of our networked experiment, i.e. the losses of audio and video streams where the `segment` and the `play` object reside on two different machine connected by a 10 Mbps Ethernet. Accordingly, our experiments were conducted with differing processor loads and network load at speeds of 30 fps (speed = 1) to 300 fps (speed = 10) for both audio and video data.

8.1 Processor Loads

In this experiment we subjected both the server and the client to varying processor loads, measured in terms of processes waiting for the processor on a 120 MHz Sparc 1 processor. The results are tabulated in terms of our metrics, aggregate and consecutive losses and drifts for both media types.

Based on our experimental results we notice the following behavior of CMT.

1. In networked video play, as in local play, playout speed increases aggregate and consecutive drops. However the variation of drops appears more wide-spread.
2. In networked audio, the `segment` object is not affected by either the load or speed, and there is wide-spread variation of drops in `play` and `device` objects.

8.2 Network Loads

In this experiment we subjected the network connecting the server and the client, to varying network loads. (measured in terms of packets per second on a 10Mbps Ethernet) between 0 to 20,000 packets/second, measured by an external *sniffer*. The results are tabulated in terms of our metrics, aggregate and consecutive losses and drifts, for both media types.

Based on our experimental results we notice the following behavior of CMT.

1. As in the case of processor loads, we noticed that increasing speed results in dropping frames. Our results show that in the given load ranges, the effect of network loads resulted in wide variations in drop rates.
2. In the audio `segment`, independent of the load range, increased speed result in more audio drops, leading to substantial packet drops at the remote `play` object. This was due to the lateness of audio frames.
3. The drifts decreased monotonically, and independently of network loads, indicating that the frames that were present did appear without much jitter, although as drops indicate many of them were not displayed.

9 Conclusions

We presented a general methodology of *user centric* or *top down* methodology of QoS study. As a case study of our methodology, we choose to characterize continuity metrics for lossy continuous media, and validate them be a user study. As the next step of the study we benchmarked CMT, a best effort test bed.




Figure 4: Video in Local Playback

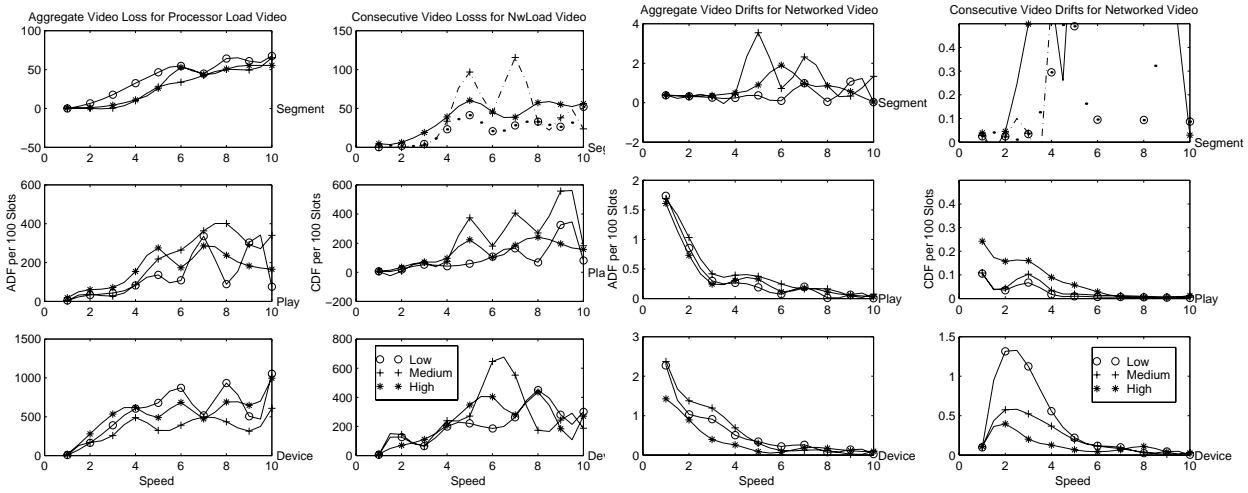


Figure 5: Video in Networked Playback

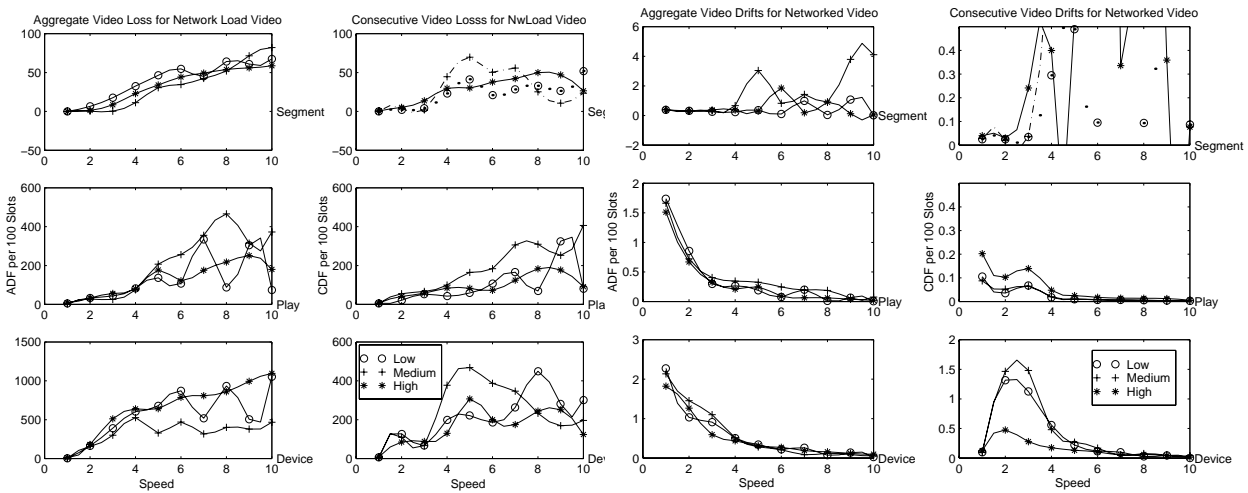


Figure 6: Video in Remote Playback against Network Loads




Figure 7: The General Trend: Seen from Local Video Segments

We have presented results of a performance evaluation of media losses in audio and video streams in CMT with respect to a set of metrics designed to measure continuity losses and timing drifts. The general trend of our results indicate that, as expected, at higher speeds, CMT loses more frames. In addition, loads increase aggregate losses at lower speeds and consecutive losses at higher speeds. It was also observed that at higher speeds, at the cost of losing more frames, the ones that are retained have smaller timing drifts. This general trend is clear for the local **segment**, as given in Figure 7.

The primary objective for this experiment was to motivate the need for QoS provisioning in CM streams and synchronization between groups of such CM streams. CMT is an ideal environment to experiment with and implement such mechanisms. Towards achieving this goal, the effect of media losses in CMT on audio-video synchronization has also been investigated in [WPV⁺98].

Our future work include making frame drops in CMT based on user specified QoS parameters. Towards this end, we are replacing appropriate object in CMT with corresponding objects that have QoS based drop policies encoded in them. Furthermore, in order to provide proper QoS provisioning at play objects, it is necessary for clients (**play** objects) to have a feed back mechanism toward the server (**segment**) objects.

References

[Ame78] American National Standards Institute (ANSI) and American Society for Quality Control (ASQC). *ANSI/ASQC Standard: Quality Term Terminology A3-1978*, a3 edition, 1978.

[Cro79] P. B. Crosby, editor. *Quality is Free*. McGraw Hill Book Company, New York, 1 edition, 1979.

[Deu] Deutsches Insitute fuer Normung Teil 11, Beuth-Varlag, Berlin. *DIN-53350*.

[Eur81] European Organization for Quality Control. *Glossary of Trems Used in Quality Control*, 5 edition, 1981.

[Jur74] J. Juran, editor. *Quality Control Handbook*. McGraw Hill Book Company, New York, 3 edition, 1974.

[MPR97] Ketan Mayer-Patel and Lawrence A. Rowe. Design and Performance of the Berkeley Continuous Media Toolkit. In Martin Freeman, Paul Jardetzky, and Harrick Vin, editors, *Proceedings of Multimedia Computing and Networking*, pages 194–206, 1997.

[MVCC92] Alejandra M. Munoz, Gail Vance Cilille, and B. Thomas Carr. *Sensory Evaluation in Quality Control*. Van Nostrand Reinhold, 1992.

[SB96] Ralf Steinmetz and Gerold Blakowski. A media synchronization survey: Reference model, specification and case studies. *IEEE Journal on Selected Areas in Communication*, 14(1):5–35, 1996.

[Seg81] H. D. Seghezzi. What is quality conformance with requirements or fitness for the intended use. *EOQC Journal*, 4:3, 1981.

[Smi] Brian Smith. Cyclic-UDP: A Priority Driven Best-Effort Protocol. http://www.cs.cornell.edu/Info/Faculty/Brian_Smith.html/Publications.

[SRS93] Brian Smith, Lawrence A. Rowe, and Yen S. Tcl distributed programming. In *Proceedings of the 1993 Tcl/Tk Workshop*, 1993.

[SRY93] Brian Smith, Larry Rowe, and S Yen. A Tcl/Tk Continuous Media Player. In *Proceedings Tcl 1993 Workshop*, June 1993.

[SW85] M. N. Sinha and W. Willborn, editors. *The Management of Quality Assurance*. John Wiley and Sons, New York, 1985.

[Swi81] Swiss Standards Association. *QS-Norm Draft*, 1 edition, 1981.

[Tow93] Don Towsley. Providing quality of service in packet switched networks. In L. Donatiello and R. Nelson, editors, *Performance Evaluation of Computer Communication Systems*, pages 560–586. Springer Verlag, 1993.

[VKvBG95] Andreas Vogel, Brigitte Kerherve, Gregor von Bochmann, and Jan Gecsei. Distributed Multimedia and QoS: A Survey. *IEEE Multimedia Magazine*, 3:10–19, Summer 1995.

[WPV⁺98] Duminda Wijesekera, Shwetal Parikh, Srivatsan Varadarajan, Jaideep Srivastava, Anil Nerode, and Mark Foresti. Performance Evaluation of Synchronization Losses in the Continuous Media Toolkit. In Harik Vin, editor, *submitted to IEEE Multimedia - 1998*, volume 2, University of Texas, Austin, October 1998.

[WS96] Duminda Wijesekera and Jaideep Srivastava. Quality of Service (QoS) Metrics for Continuous Media. *Multimedia Tools and Applications*, 3(1):127–166, September 1996.

[WSNF97] Duminda Wijesekera, Jaideep Srivastava, Anil Nerode, and Mark Foresti. Experimental Evaluation of Loss Perception in Continuous Media. *submitted to ACM Multimedia*, July 1997.