

Adaptive Internet Interactive Team Video

Dan Phung¹ Giuseppe Valetto^{1,2} Gail Kaiser¹

¹Computer Science Department ²Telecom Italia Lab
Columbia University Turin, Italy
New York City, New York

{phung, valetto, kaiser}@cs.columbia.edu

Abstract

The increasing popularity of distance learning and online courses has highlighted the lack of collaborative tools for student groups. In addition, the introduction of lecture videos into the online curriculum has drawn attention to the disparity in the network resources used by students. We present an e-Learning architecture and adaptation model called AI²TV (Adaptive Internet Interactive Team Video), a system that allows borderless, virtual students, possibly some or all disadvantaged in network resources, to collaboratively view a video in synchrony. AI²TV upholds the invariant that each student will view semantically equivalent content at all times. Video player actions, like play, pause and stop, can be initiated by any of the students and the results of those actions are seen by all the other students. These features allow group members to review a lecture video in tandem to facilitate the learning process. We show in experimental trials that our system can successfully synchronize video for distributed students while, at the same time, optimizing the video quality given actual (fluctuating) bandwidth by adaptively adjusting the quality level for each student.

Keywords: E-Learning Architectures, Borderless Education, Virtual University

1 Introduction

Life-Long and distance learning programs such as the Columbia Video Network¹ and the Stanford Center for Professional Development have evolved from FedEx-ing lecture video tapes to their off-campus students to instead streaming the videos over the Web. The lectures might be delivered “live”, but are more frequently post-processed and packaged for students to watch (and re-watch) at their convenience. This introduces the possibility of forming “study groups” among students who can view the lecture videos together and pause, rewind, or fast-forward the video to discussion points, thus approximating the pedagogically valuable discussions that occur during on-campus lectures. To that end, we provide an e-Learning architecture that can support those virtual student groups.

Using conventional Internet-video technology, *collaborative video viewing* by multiple geographically dispersed users is not yet supported. It is particularly challenging to support WISIWYS (what I see is what you see) when some of the users are relatively disadvantaged with respect to bandwidth (e.g., dial-up modems) and local computer resources (e.g., old graphics cards, small disks). We have adopted technology (developed by others, Liu and Kender [?]) for “semantically compressing” standard MPEG videos into sequences of still JPEG images. This technology automatically selects the most semantically meaningful frames to show for each time epoch, and can generate different sequences of JPEG images for a range of different compression (bandwidth) levels. This technology was designed with typical lecture videos in mind, so this approach works well for taped classes; for instance, it recognizes that it is semantically more important to see the blackboard after the instructor has finished writing, than showing the instructor’s back as he writes the same content on the board.

The remaining technical challenges are *synchronizing* and *adapting* the downloading and display of the image sequences among each of the distributed students, including support for shared video player actions such as pause. Further,

¹CVN caters to non-degree continuing education students.

if student groups do indeed sometimes pause the videos, or rewind to a point already available in local buffers (caches), it is desirable to take advantage of the then-idle network connection to prefetch future images at a higher quality level.

We have developed an approach that achieves these goals using three mechanisms working in tandem. First, the software clocks of the video clients for each student are synchronized using NTP [?]. This time is used for reference within the image sequences, where each image is associated with its start and end times relative to the beginning of the sequence. Second, the video clients communicate with each other over a distributed publish-subscribe event bus, which propagates video actions taken by one user in the group to all the other users in the group. Thus any user can select a video action, not just a “leader”.

We are particularly concerned about disenfranchised user communities that have relatively low-bandwidth over dialup, so the third and main innovation of our research concerns optimizing the video quality according to the bandwidth constraints of each user. A distributed feedback control loop dynamically adjusts each video client’s choice of both the next image to display and also the next image to retrieve from the semantic compression levels available. The controller relies on sensors embedded in each client to periodically check what image is currently displaying, whether this image is “correct” for the current NTP time compared to what other clients are viewing, which images have already been buffered (cached) at that client, and what is the actual bandwidth recently perceived at that client. Actuators are inserted into the video clients to automatically tune local configuration parameters. The controller utilizes detailed information about the image sequences available at the video server, including image start and stop times (both the individual images and their start and stop times tend to be different at different compression levels), but unlike local client data, video server data is unlikely to change while the video is showing. A single controller is used for all clients in the same student group, so it can detect “skew” across multiple clients, and may reside on the video server or on another host on the Internet.

In the next section, we further motivate the collaborative video viewing prob-

lem, provide background on the semantically compressed video repository, and explain the technical difficulties of optimizing quality while synchronizing such semantically compressed videos. The following section presents our architecture and dynamic adaptation model, and its implementation in AI²TV (Adaptive Interactive Internet Team Video). In the evaluation section, we describe the criteria used to evaluate the effectiveness of our approach, and show empirical results obtained when applied to real lecture videos distributed for a recent Columbia Video Network course. We conclude with a comparison to related work and by summarizing our contributions.

2 Motivation and Background

Correspondence courses have been available to working adult and/or geographically remote learners for over a century, e.g., the American School in Illinois ² has offered high school courses since 1897, and the University of Wyoming^[?] began offering extension courses in 1892, launching a full-fledged college distance learning program in 1906. Correspondence courses have traditionally been designed for individual students with a self-motivated learning style, studying primarily from text materials.

An NSF Report ^[?] discusses how technology, from radio to television, to audio and video cassettes, to audio and video conferencing, has affected distance education. These technologies have enabled educational institutions to offer certification and degree tracks using live or pre-taped audio and/or video of regular on-campus classroom lectures. The report states that the recent use of Internet technologies, especially the Web, has “allowed both synchronous and asynchronous communication among students and between faculty and students” and has “stimulated renewed interest in distance education”. It also mentions that “stimulating interaction among students” can help reduce dropout rates, which it says may be higher in distance education than in traditional courses. Finally, it cites some studies that “suggest the Web is superior to earlier distance education

²<http://www.americanschoolofcorr.com/>

technologies because it allows teachers to build collaborative and team-oriented communities rather than either the passive classes of the conventional academy or the individual study of traditional correspondence courses”.

Today’s equivalent of correspondence courses are often offered online through a Web portal interface, with some for-profit schools like University of Phoenix³ which are primarily online. Some tools such as instant messaging, application and desktop sharing, and co-browsing [?] can facilitate the communicative aspects of synchronous collaboration, but are not designed specifically for educational purposes. Support for synchronous collaboration remains a major concern in courses where group work is encouraged [?], yet there are few educational tools that allow synchronous collaboration across a group of online students [?]. However, it seems that Web-based video streaming should enable synchronous collaboration “situated” by collaborative lecture video viewing, approximating the experience of on-campus students physically attending the lecture and class discussion.

Our AI²TV project aims to contribute to the area of synchronous collaboration support for life-long and distance education, specifically in the context of collaborative video viewing, to foster virtual classrooms and borderless education. Our design is intended for small classes or small study groups within a larger class, and reaches out to disenfranchised user communities with dial-up level bandwidths, who still constitute a significant portion of the Internet user community [?], to allow them to collaborate with other users that enjoy broadband or higher bandwidth resources.

Viewing video on the Internet usually requires relatively high bandwidth resources, and low-bandwidth or lossy network connections can lead to lost video content. This is particularly a problem for group review of lecture videos, if different members of the group miss different portions of the video or fall behind to different degrees due to extensive buffering. Furthermore, disadvantages in network and computing resources may make it difficult if not impossible – with current Internet video technology – for some students to participate in

³<http://www.phoenix.edu/>

collaborative lecture video viewing at all.

Technically, collaborative video viewing poses a twofold problem: on the one hand, it is mandatory to keep all users synchronized with respect to the content they are supposed to see at any moment during play time; on the other hand, it is important to provide each individual user with a level of quality that is optimized with respect to the user's available resources, which may vary during the course of the video.

One solution to the problem of balancing the group synchronization requirement with the optimization of individual viewing experiences is to use videos with cumulative layering [?], also known as scalable coding [?]. In this approach, the client video player selects a quality level appropriate for that client's resources from a hierarchy of several different encodings for that video. Thus a client could receive an appropriate quality of video content while staying in sync with the other members of the group.

We use *semantic compression* to produce a video with cumulative layering. The semantic compression algorithm developed by Liu and Kender [?] reduces a video to a set of semantically significant key frames. That tool operates on conventional MPEG format videos and outputs sequences of JPEG frames, some of which are displayed in figure 1. The semantic compression algorithm profiles video frames within a sliding time window and selects key frames that have the most semantic information with respect to that window. By increasing the size of the window, a key frame will represent a larger time slice, which means that a larger window size will produce less key frames as compared to a smaller window size setting.

A conceptual diagram of a layered video produced from this semantic compression is shown in figure 1. Note that the semantic compression algorithm produces an effectively random distribution of key frames, hence the video produced by the package plays back at a *variable* frame rate. The variability in the frame rate is most significant when there are pockets of relatively high frequency semantic change, which result in sections in the video that demand a higher frame rate. The variable frame rate video adds substantial complexity to

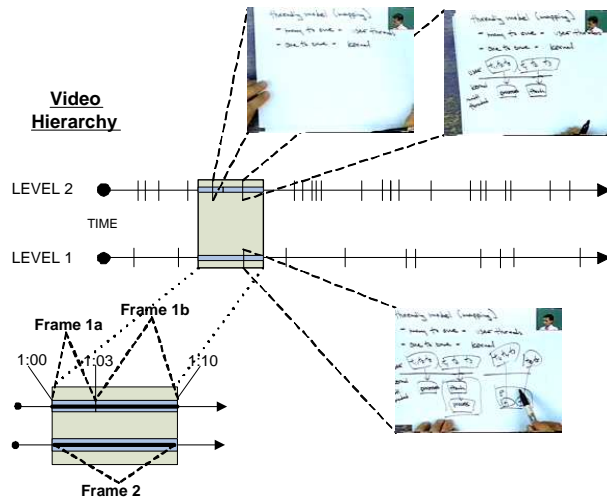


Figure 1: Semantic Video Scenario

the bandwidth demands of the client.

In figure 1, the bottom-left in-set shows the juxtaposition of individual frames from two different quality levels. Each frame has a representative time interval $[\text{start}:\text{end}]$. For the higher level, Frame 1a represents the interval from 1:00 to 1:03, and Frame 1b represents the interval from 1:04 to 1:10. For the lower level, Frame 2 represents the entire interval from 1:00 to 1:10. In this diagram, Frame 2 is semantically equivalent to Frame 1a and 1b together. However, in real JPEG frame sequences produced from the same MPEG video for different quality levels, the start and end times of frame sets rarely match up as ideally as in our example.

Through the use of the Liu/Kender video compression algorithm, we can potentially provide semantically equivalent content to a group of students with diverse resources by adjusting the compression level assigned to each client while the students are watching the video. Thus for our purposes, synchronization of collaborative video boils down to showing semantically equivalent frames at all times.

To adjust the video clients in response to the changing environment, we use

an “autonomic” controller⁴, to maintain the synchronization of the group of video clients while simultaneously fine tuning the video quality seen by each student. The controller remains conceptually separate from the controlled video system, employing our decentralized workflow engine, named Workflakes [?]. Said workflow coordinates the behavior of software entities, as opposed to conventional human-oriented workflow systems; the use of workflow technology for the specification and enactment of the processes coordinating software entities was previously suggested by Wise et al. [?]. Workflakes has previously been used in a variety of more conventional “autonomic computing” domains, where it orchestrates the work of software actuators to achieve the fully automated dynamic adaptation of distributed applications [?, ?, ?]. In the context of AI²TV, Workflakes monitors the video clients and consequently coordinates the dynamic adjustment of the compression (quality) level currently assigned to each client.

3 Architecture and Adaptation Model

3.1 System Architecture

AI²TV involves several major components: a video server, video clients, an autonomic controller, and a common communications infrastructure, as shown in figure 2.

The video server provides the educational video content to the video clients for viewing. Each lecture video is stored in the form of a hierarchy of versions, produced by running the semantic compression tool multiple times with settings for different compression levels. Each run produces a sequence of JPEG frames with a corresponding frame index file. The task of the video server is simply to provide remote download access to the collection of index files and frames over HTTP.

The task of each video client is to acquire video frames, display them at the correct times, and provide a set of basic video functions. Taking a func-

⁴The term *autonomic* is borrowed from IBM to mean a self-managing system that uses a (software) feedback control loop

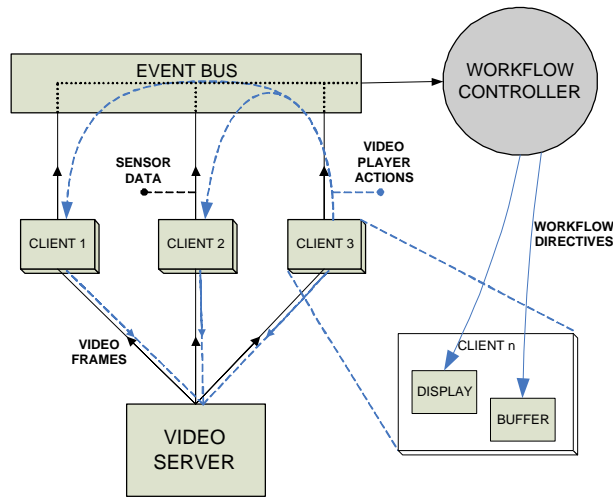


Figure 2: AI²TV Architecture

tional design perspective, the client is composed of four major modules: a time controller, video display, video buffer that feeds the display, and a manager for fetching frames into the buffer.

The time controller’s task is to ensure that a common video clock is maintained across clients. It relies on NTP to synchronize the system’s software clocks, therefore ensuring a common time base from which each client can reference the video indices. Using this foundation, the task of each client is simplified to displaying the client’s needed frame at the correct time. Since all the clients refer to the same time base, then all the clients are showing semantically equivalent frames from the same or different quality levels.

The video display renders the JPEG frames at the correct time into a window and provides a user interface for play, pause, goto and stop. When any participant initiates such an action, all other group members receive the same command, thus all the video actions are synchronized. Video actions are time stamped so that clients can respond to those commands in reference to the common time base. The video display knows which frame to display by using the current video time and display quality level to index into the frame index for the representative frame. Before trying to render the frame, it asks the video

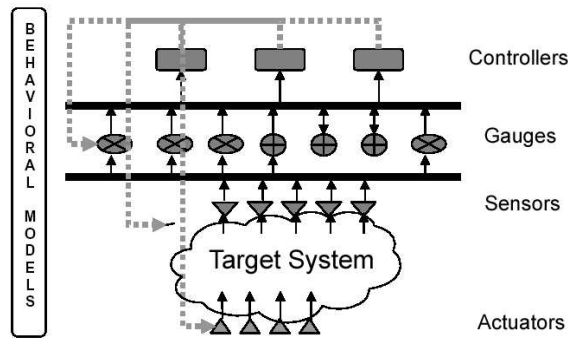


Figure 3: Conceptual Reference Architecture

buffer manager if the needed frame is available. The video display also includes a control hook that enables external entities, like the autonomic controller, to adjust the current display quality level.

The video manager constitutes a downloading daemon that continuously downloads frames at a certain level into the video buffer. It keeps a hash of the available frames and a count of the current reserve frames (frames buffered) for each quality level. The buffer manager also includes a control hook that enables external entities to adjust the current downloading quality level.

The purpose of the autonomic controller is to ensure that, given the synchronization constraint, each client plays at its highest attainable quality level. The controller is itself a distributed system, whose design derives from a conceptual reference architecture for autonomic computing platforms proposed by Kaiser *et al.* [?], which is shown in figure 3.1. The architecture provides an end-to-end closed control loop, in which sensors attached to a generic (possibly legacy) target system continuously collect and send streams of data to gauges. The gauges analyze the incoming data streams and recognize adverse conditions that need adaptation, relaying that information to controllers. The controllers coordinate the expression and orchestration of the workflow needed to carry out the adaptation. At the end of the loop, actuators attached to the target system effect the needed adjustments under the supervision of the controller.

In the AI²TV case, sensors at each client monitor for currently displayed

frame, its quality level, the quality level currently being fetched by the manager, the time range covered by buffer reserve frames, and the current bandwidth. Gauges are embedded together with the controller for expediency in design and to minimize communication latency. They receive the sensor reports from individual clients, collect them in buckets, similar to the approach in [?], and pass the bucket data structure to the controller’s coordination engine. A set of helper functions tailored specifically for this application operate on this data structure and produce triggers for the coordinator. When a trigger is raised, the coordination engine enacts an adaptation scheme, basically a workflow plan, which is executed on the end hosts by taking advantage of the hooks provided to the actuators by the clients.

Communication among the video clients, as well as between the sensors and actuators at the clients and the autonomic controller, is provided by a publish-subscribe event bus. There are three kinds of events: video player actions, sensor reports, and adaptation directives (see figure 2).

3.2 Adaptation Model

The adaptation scheme consists of two levels: a higher level data flow, and a lower level adjustment heuristic. The former directs the flow of data through a logical sequence to provide a formal decision process, while the latter provides the criteria as to when to make certain adjustments.

The higher level logic is shown in figure 4. The diagram shows the task decomposition hierarchy according to which the adaptation workflow unfolds. Note that the evaluation of clients’ state with respect to the group (`EvaluateClient`) and the issuing of adaptation directives (`AdaptClient`) is carried out as a set of parallel steps. Also note that the multiplicity of those parallel steps is dynamically determined via the number of entries in the `clients` variable, which maps to a collection of AI²TV clients.

The adaptation scheme at the lower level falls into two categories: directives that adjust the client in response to relatively low bandwidth situations, and those that take advantage of relatively high bandwidth situations.




Figure 4: AI²TV Workflow diagram

In the situation where a client has relatively low bandwidth, the client may not be able download the next frame at the current quality level by the time it needs to begin displaying that frame. Then both the client and buffer quality levels are adjusted downwards one level. If the client is already at the lowest level (among those available from the video server), the controller will calculate the next possible frame that most likely can be successfully retrieved before its own start time while remaining synchronized with the rest of the group. The client will then be directed to jump ahead to that frame.

To take advantage of relatively high bandwidth situations, the buffer manager will start to accumulate a reserve buffer. Once the buffer reaches a threshold value (e.g., 10 buffered frames), the controller will direct the manager to start fetching frames at a higher quality level. Once sufficient reserve is accumulated also at that higher level, the client is then ordered to display frames at that quality level. If the bandwidth drops before the buffer manager can accumulate enough frames in the higher-level reserve, the buffer manager is dropped back down one quality level.

3.3 Implementation

Our system is implemented in Java. The video client uses `javax.swing` to render JPEG images. The controller, Workflakes, is built on top of the open-source Cougaar multi-agent system⁵, which it extends to allow the orchestration of distributed software agents for autonomic purposes (explained further in [?]). We used the Little-JIL graphical workflow specification language [?] for defining adaptation plans. We chose a freely available, content-based, publish-subscribe event system, Siena [?], as our communication bus.

4 Evaluation

Our assessment considers the ability of AI²TV to synchronize the clients and to optimally adjust their video quality. Our results were computed from client configurations simulating small study groups which consisted of 1, 2, 3, and 5 clients together running a semantically summarized video for 5 minutes, with sensors probing clients state every 5 seconds. The compression hierarchy we employed has 5 different quality levels.

We define a baseline client against which the performance of our approach can be compared. The baseline client's quality level is set at the beginning of the video and not changed thereafter, using a value we identify as the average bandwidth per level. This value is computed by summing the total size in bytes of all frames produced at a certain compression level and dividing by the total video time. This value provides the bandwidth needed, on average, for the buffer manager to download the next frame on time. We provide the baseline client with the corresponding bandwidth for its chosen level by using a bandwidth throttling tool `shaperd`⁶ to adjust the bandwidth between that client and the video server. Note that using the average as the baseline does not account for the inherent variability in video frame rate and likely fluctuations in real-world network bandwidth, where adaptive control can make a difference.

⁵<http://www.cougaar.org/>

⁶<http://freshmeat.net/projects/shaperd/>

Each controller-assisted client is assigned an initial level in the compression hierarchy and the same bandwidth as the baseline client for that hierarchy level. For each experimental trial, we record any differences resulting from the controller's adaptation of the clients' behavior *versus* the behavior of the baseline client, with respect to synchrony and frame rate.

4.1 Evaluating Synchronization

The primary goal of our system is to provide synchronous viewing of lecture videos to small groups of geographically dispersed students, some possibly with relatively meager resources. Our initial experiments evaluate the level of synchronization for several small groups of clients, where each group is involved in a video session. Each client is preset at a designated level of compression and given the average baseline bandwidth required for that compression level. To measure the effectiveness of the synchronization, we probe the video clients at periodic time intervals and log the frame currently being displayed. This procedure effectively takes a series of system snapshots, which we can evaluate for synchronization correctness. We check whether the frame being displayed at a certain time corresponds to one of the valid frames for that time, on *any* quality level. We allow an arbitrary level here because the semantic compression algorithm ensures that all frames designated for a given time will contain semantically equivalent information. We obtain a score by summing the number of clients not showing an acceptable frame and normalizing over the total number of clients. A score of 0 indicates a fully synchronized system.

These experiments showed a total score of 0 for all trials, meaning that all of the clients were viewing appropriate frames when probed. Notwithstanding the variations in the frame rate and/or occasional fluctuations in the actual bandwidth of the clients, no frames were missed. This result demonstrates that the chosen baseline combinations of compression levels and throttled bandwidths do not push the clients beyond their bandwidth resource capacity.

Then we ran another set of experiments, in which the clients were assigned more casually selected levels of starting bandwidths. Said casual selection is rep-

representative of real-world situations, like receiving Internet audio or audio/video streams, where users must choose a desired frame rate for the transmission of the content. The user may have been informed that she is allocated a certain bandwidth level from her Internet service provider, but may actually be receiving a significantly lower rate. The clients were assigned bandwidths one level lower than the preset quality level. We ran this set of experiments first without the aid of the autonomic controller and then with it. In the former case, clients with insufficient bandwidth were stuck at the compression level originally selected, and thus missed an average of 63% of the needed frames. In the latter case, the same clients only missed 35% of the needed frames. Although both situations show a significant amount of missed frames, these results provide evidence of the benefits of the adaptive scheme implemented by the autonomic controller.

4.2 Evaluating Quality of Service

The most interesting technical innovation of the AI²TV system is our autonomic controller approach to optimizing video quality. Here we analogously use a scoring system relative to the baseline client's quality level. We give a weighted score for each level above or below the baseline quality level. The weighted score is calculated as the ratio of the frame rate of the two levels. For example, if a client is able to play at one level higher than the baseline, and the baseline plays at an average n frames per second (fps) while the level higher plays at $2*n$ fps, the score for playing at the higher level is 2. The weighted score is calculated between the computed average frame rates of the chosen quality levels. Theoretically, the baseline client should receive a score of 1. Note that we formulated this scoring system because other scoring systems (e.g., [?, ?, ?]) measure unrelated factors such as the synchronization between different streams (audio and video), image resolution, or human perceived quality, and are not constrained by the group synchronization requirement. This restriction mandates a scoring system sensitive to the relative differences between quality hierarchies.

Our experiments show that baseline clients scored a group score of 1 (as expected) while the controller-assisted clients scored a group score of 1.25. The

one-tailed t-score of this difference is 3.01, which is significant for an α value of .005 (N=17). This result demonstrates that using the autonomic controller enabled our system to achieve a significant positive difference in the quality of service (QoS) aspect that relates to received frame rate. Note that the t-score does not measure the degree of the positive difference: To demonstrate the degree of benefit, we measure the proportion of additional frames that each client is able to enjoy. We found that, overall, those clients received 20.4% (\pm 9.7, N=17) more frames than clients operating at a baseline rate.

Running the client close to or at a level higher than the average bandwidth needed puts the client at risk for missing more frames, because the autonomic controller is trying to push the client to a better but more resource-demanding level. To measure whether the controller-assisted client is adversely exposed to a higher risk of missing frames, we also count the number of missed frames during a video session. The scoring is a simple count of the missed frames. Note that this scoring is kept separate from the measure of the relative quality to discriminate between levels of concern, although they both indicate QoS characteristics.

In all of our experiments, there was only one instance in which a controller-assisted client missed some frames, in particular two consecutive frames. Upon closer inspection, the time region during that event showed that the semantically compressed video demanded a higher frame rate at the same time that the network bandwidth available to that client was relatively low. The client was able to consistently maintain a high video quality level after this epoch.

Our AI²TV system can achieve collaborative video viewing using relatively naive NTP-based synchronization, without the autonomic controller. But in typical real-world scenarios, network bandwidth varies over time, plus the variable frame rate of semantically compressed video does not permit the client to make an informed decision about the most appropriate quality level for the next frames without adaptive technology akin to our controller. Our experimental data shows that the autonomic controller makes a significant positive difference in achieving higher QoS.

5 Related Work

Yin *et al.* [?] provide a system for an adaptive distributed multimedia system that is based on streaming, multicast and compression technology. They show that they can provide an improved level of QoS, but do not discuss user-level action synchronization, and use quality degradation rather than semantic compression to adapt to client resource constraints. Cen *et al.* provide a distributed real-time MPEG player that uses a software feedback loop between a single server and a single client to adjust frame rates [?]. Their architecture incorporates feedback logic within each video player and does not support synchronization across a group of players, while the work presented here explicitly supports the synchronization of semantically equivalent video frames across a small group of clients.

An earlier implementation of AI²TV is described in [?]. In that version, a collaborative virtual environment (CVE) supported a variety of team interactions [?], with the optional lecture video display embedded in the wall of a CVE “room”. Video synchronization data was piggybacked on top of the UDP peer-to-peer communication used primarily for CVE updates, which did not work very well due to the heavy-weight CVE burden on local resources.

Our approach to synchronization can be classified as a distributed adaptive scheme that employs a global clock and operates in a proactive way. The most significant difference compared to other approaches, such as the Adaptive Synchronization Protocol [?], the work of Gonzalez and Adbel-Wahab [?], or that of Liu and El Zarki[?], is that our approach is not based on the idea of play-out delay. Instead, we take advantage of layered semantic compression coupled with buffering to “buy more time” for clients that might not otherwise be able to remain in sync, by putting them on a less demanding level of the compression hierarchy.

Liu *et al.* provide a comprehensive summary of the mechanisms used in video multicast for quality and fairness adaptation as well as network and coding requirements [?]. To frame our work in that context, our current design and

implementation models a single-rate server adaptation scheme to each of the clients because the video quality we provide is tailored specifically to that client's network resources. The focus in our work is directed towards the client-side end-user perceived quality and synchrony, so we did not utilize the most efficient server model. The authors believe that it would be trivial to substitute in a simulcast server adaptation model [?].

6 Conclusion

We present an e-Learning architecture and prototype system that allows small, geographically dispersed student groups to collaboratively view lecture videos in synchrony. To accommodate disenfranchised users with relatively low-bandwidth, AI²TV employs an “autonomic” (feedback loop) controller to dynamically adapt the video quality according to each client's network resources. We rely on a semantic compression algorithm to guarantee that the semantic composition of the simultaneously viewed video frames is equivalent for all clients. Our system distributes appropriate quality levels of video to clients, and automatically adjusts them according to their current bandwidth resources. We have demonstrated the advantages of this approach through experimental trials using bandwidth throttling to show that our system can provide synchronization of video together with optimized video quality to distributed student groups.

7 Acknowledgments

We would like to thank John Kender, Tiecheng Liu, and other members of the High-Level Vision Lab for their assistance in using their lecture-video semantic compression software. We would also like to thank other members of the Programming Systems Lab, particularly Matias Pelenur who implemented PSL's Little-JIL interpreter on top of Workflakes/Cougaar and Suhit Gupta for invaluable commentary and ideas. Little-JIL was developed by Lee Osterweil's LASER lab at the University of Massachusetts, Amherst. Cougaar was developed by a DARPA-funded consortium; our main Cougaar contact was Nathan Combs of BBN. Siena was developed by the University of Colorado, Boulder, in Alex Wolf's SERL lab. PSL is funded in part by National Science Foundation grants CNS-0426623, CCR-0203876 and EIA-0202063, and by Microsoft Research.

References

- [1] The Application and Implications of Information Technologies in Postsecondary Distance Education: An Initial Bibliography. Technical Report NSF 03-305, National Science Foundation, Division of Science Resources Statistics, December 2002.
- [2] S. Baqai, M. F. Khan, M. Woo, S. Shinkai, A. A. Khokhar, and A. Ghafoor. Quality-based evaluation of multimedia synchronization protocols for distributed multimedia information systems. *IEEE Journal of Selected Areas in Communications*, 14(7):1388–1403, 1996.
- [3] L. A. Burgess and S. D. Strong. Trends in online education: Case study at southwest missouri state university. *Journal of Industrial Teacher Education*, 19(3), 2003.
- [4] M. Capps, B. Laddi, D. Stotts, and L. Nyland. Educational applications of multi-client synchronization through improved web graph semantics. In *5th International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 1996.
- [5] A. Carzaniga, D. S. Rosenblum, and A. L. Wolf. Design and evaluation of a wide-area event notification service. *ACM Transactions on Computer Systems*, 19(3):332–383, Aug. 2001.
- [6] A. G. Cass, Barbara Staudt Lerner, E. K. McCall, L. J. Osterweil, Stanley M. Sutton, Jr., and A. Wise. Little-JIL/Juliette: A Process Definition Language and Interpreter. In *22nd International Conference on Software Engineering*, pages 754–757, June 2000.
- [7] A. L. Corte, A. Lombardo, S. Palazzo, and G. Schembra. Control of perceived quality of service in multimedia retrieval services: Prediction-based mechanism vs. compensation buffers. *Multimedia Systems*, 6(2):102–112, 1998.
- [8] S. E. Dossick and G. E. Kaiser. CHIME: A Metadata-Based Distributed Software Development Environment. In *Joint 7th European Software Engineering Conference and 7th ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, pages 464–475, 1999.
- [9] L. Gautier and C. Diot. Design and evaluation of mimaze, a multi-player game on the internet. In *International Conference on Multimedia Computing and Systems*, pages 233–236, 1998.
- [10] A. J. Gonzalez and H. Adbel-Wahab. Lightweight stream synchronization framework for multimedia collaborative applications. In *5th IEEE Symposium on Computers and Communications*, July 2000.
- [11] S. Gupta and G. Kaiser. A Virtual Environment for Collaborative Distance Learning With Video Synchronization. In *7th IASTED International Conference on Computers and Advanced Technology in Education*, August 2004.
- [12] G. Kaiser, P. Gross, G. Kc, J. Parekh, and G. Valetto. An Approach to Autonomizing Legacy Systems, in *Workshop on Self-Healing, Adaptive and Self-Managed Systems*, in *Workshop on Self-Healing, Adaptive and Self-Managed Systems*, June 2002.
- [13] G. Kaiser, J. Parekh, P. Gross, and G. Valetto. Kinesthetics eXtreme: An External Infrastructure for Monitoring Distributed Legacy Systems. In *5th Annual International Active Middleware Workshop*, June 2003.

- [14] G. Kaiser, J. Parekh, P. Gross, and G. Valetto. Retrofitting Autonomic Capabilities onto Legacy Systems. Technical Report CUCS-026-03, Columbia University Department of Computer Science, October 2003.
- [15] B. S. Lemer, E. K. McCall, A. Wise, A. G. Cass, L. J. Osterweil, and S. M. S. Jr. Using little-jil to coordinate agents in software engineering. In *Automated Software Engineering Conference*, September 2000.
- [16] W. Li. Overview of the fine granularity scalability in mpeg-4 video standard. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(3):301–317, March 2001.
- [17] X. Li, M. H. Ammar, and S. Paul. Video multicast over the internet, April 1999.
- [18] H. Liu and M. E. Zarki. A synchronization control scheme for real-time streaming multimedia applications. In *Packet Video 2003*, April 2003.
- [19] J. Liu, B. Li, and Y.-Q. Zhang. Adaptive video multicast over the internet. *IEEE Multimedia*, 10(1):22–33, January/March 2003.
- [20] T. Liu and J. R. Kender. Time-constrained dynamic semantic compression for video indexing and interactive searching. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 531–538, 2001.
- [21] S. McCanne, V. Jacobson, and M. Vetterli. Receiver-driven layered multicast. In *ACM SIGCOMM*, volume 26,4, pages 117–130, New York, AUG 1996. ACM Press.
- [22] J. Miller, C. Ditzler, and J. Lamb. Reviving a Print-based Correspondence Study Program In the Wake of Online Education. In *American Association for Collegiate Independent Study: Distance Learning: Pioneering the Future*, October 2003.
- [23] D. L. Mills. Network time protocol. RFC 958, 1985.
- [24] M. Richtel. In a Fast-Moving Web World, Some Prefer the Dial-Up Lane. *The New York Times*, April 19 2004.
- [25] K. Rothermel and T. Helbig. An Adaptive Protocol for Synchronizing Media Streams. *Multimedia Systems*, 5:324–336, 1997.
- [26] G. Valetto. *Orchestrating the Dynamic Adaptation of Distributed Software with Process Technology*. PhD thesis, Columbia University, April 2004.
- [27] G. Valetto and G. Kaiser. Using Process Technology to Control and Coordinate Software Adaptation. In *International Conference on Software Engineering*, May 2003.
- [28] J. Walpole, R. Koster, S. Cen, C. Cowan, D. Maier, D. McNamee, C. Pu, D. Steere, and L. Yu. A Player for Adaptive MPEG Video Streaming Over The Internet. In *26th Applied Imagery Pattern Recognition Workshop*. SPIE, October 1997.
- [29] Y. Wang, J. Ostermann, and Y.-Q. Zhang. *Video Processing and Communications*. Prentice Hall, 2002.
- [30] J. G. Wells. Effects of an on-line computer-mediated communication course. *Journal of Industrial Technology*, 37(3), 2000.
- [31] H. Yin, C. Lin, J.-J. Zhuang, and Q. Ni. An adaptive distance learning system based on media streaming. In *International Conference on Web-Based Learning*, pages 184–192, 2004.