Experiences in Personal Lecture Video Capture

Surendar Chandra

Abstract—The ability of lecture videos to capture the different modalities of a class interaction make them a good review tool. Multimedia capable devices are ubiquitous among contemporary students. Many lecturers are leveraging this popularity by distributing videos of lectures. They depend on the university to provide the video capture infrastructure. Some universities use trained videographers. Though they produce excellent videos, these efforts are expensive. Several research projects automate the video capture. However, these research prototypes are not readily deployable because of organizational constraints. Rather than waiting for the university to provide the necessary infrastructure, we show that instructors can personally capture the lecture videos using off-the-shelf components. Consumer grade high definition cameras and powerful personal computers allow instructor captured lecture videos to be as effective as the ones captured by the university. However, instructors will need to spend their own time on the various steps of the video capture workflow. They are also untrained in media capture; the capture mechanisms must be simple. Based on our experience in capturing lecture videos over three and a half years, we describe the technical challenges encountered in this endeavor. For instructors who accept the educational value of distributing lecture videos, we show that the effort required to capture and process the videos was modest. However, most existing campus storage and distribution options are unsuitable for the resource demands imposed by video distribution. We describe the strengths of several viable distribution alternatives. The instructors should work with the campus information technology personnel and design a distribution mechanism that considers the network location of the students.

Index Terms—E-learning tools, lecture notes, mobile and personal devices, nomadic learning environments.

✦

1   INTRODUCTION

EDUCATION is a lifelong endeavor with introductory courses providing the necessary foundation for more advanced topics. However, students miss some lectures. They also forget the topics that were covered in earlier lectures. Hence, they desire lecture review tools for use either within the same course or for use in the future.

Traditionally, students took notes during the lectures and then saved them for future use. When they missed lectures, they borrowed notes from their peers. However, taking written notes disrupts the student from paying full attention to the lecture. Also, incomplete notes are inadequate when the student had also forgotten the context within which a particular topic was discussed.

In the extreme, students audit the same course in the future (e.g., students audit prerequisites). However, the future class will likely be offered by another instructor with a different mixture of topics. The student might have to attend the entire course in order to relearn the relevant topics. Students prefer to review the topics using adequate lecture materials from the course that they themselves had experienced.

Many instructors assist the students by providing a copy of the lecture slides. This is especially easy when the slides were prepared for electronic presentation during the lecture. However, the slides accurately represent the lecture as it was prepared for presentation; they do not include any impromptu discussions and blackboard illustrations. Students are still required to take additional notes to complement the slides.

Recently, an audio recording of the lecture, captured either automatically or by the instructor has become a popular review tool. The audio clips are typically distributed as a podcast [1], [2]. McKinney et al. [3] report that psychology students who took notes while listening to the podcasts performed better than students who attended the actual lectures.

Some instructors [4] are also distributing a video of the screen that was projected in the classroom as a screencast. The `Record Narration` feature of Microsoft Powerpoint associates the latest audio narration with each slide; these audio annotated slides can be stored as a movie. Tools such as Camtasia and Lecturnity provide screencast capture functionality while others [4] developed their own tool. Screencasts are useful for instructors who use computers for all illustrations (e.g., using tablet PCs and pen tablets). However, other instructors prefer blackboards. Lanir et al. [5] observed that instructors employed a greater variety of instructional techniques while using blackboards than when using powerpoint. Such interactions are not captured by screencasts.

Podcasts and screencasts alone do not capture the full class interaction. Typically, instructors project a prepared set of slides. They expand on these slides using the blackboard. They also highlight aspects of the slides using mouse gestures and with laser pointers. Classes also include lively discussion between students and the instructor. Ideally, one needs to capture the data from all these modalities for an effective lecture review.

Videos can capture much of the data associated with each modality. Different mechanisms are required to capture each modality. Some events are captured using fixed mechanisms (e.g., capturing the LCD projection using the NCast Telepresenter) while capturing a video of the instructor might require object tracking technologies in

• *The author is with the FX Palo Alto Laboratory, 3400 Hillview Avenue, Palo Alto, CA 94304. E-mail: surendar@acm.org.*

order to follow the lecturer. These capture mechanisms should also be actively deployed in the lecture hall.

Video streams are large. It is not feasible to distribute each individual stream; further processing must create a single video that captures all the interactions.

Videographers and automated mechanisms dominate multiple modality lecture video capture. Both these approaches are deployed and maintained by the university with minimal effort required from the instructor.

1. *Capture using skilled videographers.* Videographers are trained to use techniques such as pan, zoom, and overlay to create a video that captures all the interesting events that occurred during a lecture.

   Some universities are already addressing the problem of students in different locations and lack of synchronous meeting times using distance education and online courses, respectively. They capture and broadcast the lectures using videographers. These universities can simply distribute these videos for review purposes.

   This approach poses two significant disadvantages:

   a. Production cost is prohibitive: the videographers need to be present during the entire lecture which adds to the cost of video capture. The in-house audio/visual department charges US$100/hr for video recording and US$120/hr for editing and digitization of the video. Each course typically meets for forty lectures; potentially costing over US$8,000 per semester. The Educational Technology Services (ETS) at Berkeley[1] provides a richer capture option. However, they charge US$535 for set up and US$572/hr to capture and distribute a video containing an audio, video, and screencast of a lecture.
   b. Videographers are not always familiar with the topics being covered: their notion of important events need not coincide with those of the instructor. For example, they might zoom in on the instructor when the contents on the blackboard were more important.

2. *Automated video capture.* Many research efforts address the videographer's expense by automating aspects of the capture workflow.

   Brotherton and Abowd [6] described their experiences in collaboratively creating lecture review notes among the instructor and the students in a fully instrumented lecture hall [7]. They observed that videos were not popular because of the poor quality capture and inadequate network resources to remotely access them. The authors recognized the value of videos especially in disambiguating pronouns from the audio track. However, the technology limitations experienced by them are no longer applicable either for capture, processing, or consumption of high quality media. Mukhopadhyay and Smith [8] developed a system that combined the

video streams from a static overview camera as well as a stream from a tracking camera along with the lecture slides to create a synchronized presentation media. These tools were used to further develop mechanisms to capture and distribute lecture videos as the Berkeley Internet Broadcasting System (BIBS) [9], [10]. Similarly, Rui et al. [11], [12], [13] developed a video capture system that fully automated the lecturer and audience tracking and performed all the capture functionality while achieving the video quality close to that of human-operated systems.

Other projects enabled search capabilities. Ziewer [14] captured the screen contents using VNC. They created a fully indexed and searchable videos using VNC protocol messages, instructor annotations and through an external optical character recognition program. Similarly, Hilbert et al. [15] automatically captured the slide projection using a specialized hardware. They used optical character recognition to segment the videos and index the various slides along with the audio narration. Müller and Ottmann [16] focused on automated authoring and retrieval of lecture videos. Repp et al. [17] automated the indexing process of stored lecture videos in order to ease content-based browsing. Adcock et al. [18] created a searchable text index of the slides from publicly available lectures videos. The performance of this system is adversely affected by video overlays and authors present mechanisms to improve recognition for these videos.

Few research projects were transitioned to a large scale production service. The production BIBS [19] service is currently available in five lecture halls and uses a mixture of custom and off-the-shell tools. Burdet et al. [20] describe the effort at the University of Geneva to automate the lecture capture. The faculty collaborated with the IT staff to automatically capture the videos. In older classrooms that were not fitted with modern A/V capture infrastructure, they developed and deployed a custom capture solution using Mac Mini computers. Their solution is actively deployed in 35 lecture halls. Eth Zurich is using the Replay[2] system to manage and distribute audiovisual recordings. Their software is available through the opencast initiative (opencastproject.org). Talkminer [18] provides slide search capability for lecture videos as a public service at talkminer.com.

Transitioning prior research and implementing an automated capture system is difficult. Researchers spent most of their effort in developing novel techniques to automate the capture rather than in developing deployable solutions. Ease of deployment and maintenance requires using commodity components that are inexpensive and easily available. Currently, we require considerable technical expertise and customized hardware and software to implement, deploy, and maintain these systems. Such resources are not available in many universities.

1.1 Our Approach: Video Capture by Instructors

For the vast majority of instructors who are not in an institution that had deployed these prior systems, video capture remains inaccessible. Instead of waiting for these universities to capture the videos, we advocate an approach where each faculty member acts as the videographer. Our approach depends on the availability of easy to use, inexpensive, and off-the-shelf components. Conversely, we could not use research technologies that are unavailable in off-the-shelf software. For example, automatic annotation and rich search capabilities [21] are not yet available in commercial video processing software. The accuracy of automatic YouTube and Talkminer annotation of our own lecture videos is currently dismal.

Our approach is feasible because of the emergence of inexpensive high definition (HD) cameras along with off-the-shelf video processing software. With proper framing, the single HD stream can capture all the modalities in the lecture with sufficient detail, obviating the need to capture and mix multiple streams. The instructional technology department currently allows faculty members to borrow wireless microphones. These services can be used to also share cameras among all the instructors.

Instead of using trained videographers, our approach shifts the entire capture burden to the instructor. Faculty members are already busy in fulfilling their pedantic duties. The primary challenge for our approach is to reduce the amount of preparation time required for setting up the equipment as well as in any postprocessing steps.

We describe our experience in capturing and distributing lectures over seven semesters. We used off-the-shelf software and hardware for this endeavor. For each lecture, we required about five minutes to set up and pack our video gear; well within the time allotted by the university between successive lectures. The videos were transferred from the camera in an hour with another thirty minutes spent for adding the video annotations. We distributed HD as well as standard definition (SD) videos and audio of the lecture. The video transcoding operation required to create variations of the video for distribution can take up to a few hours of processing on modern laptops (without instructor intervention). Students can then use these variations in a wide variety of devices; making the effort to create them worthwhile. Note that viewing the video in the limited display of a portable player will lose many of the stated advantages of HD video. Our experience shows that faculty members can perform the different steps with minimal effort.

Our students expressed similar opinions on the usefulness of our videos as with ones captured by other means. We reconfirmed prior observations that lecture videos are a useful review tool. We also did not observe any significant drop in class attendance.

The administrators viewed our effort as an attempt at improving the visibility of the university and as a recruitment tool and hence were generally supportive. However, the university was unwilling to invest significant resources in deploying automated capture mechanisms. This provided the impetus to develop our approach.

The final challenge was in the choice of a distribution mechanism. The university provided instructional storage was inadequate for video distribution. Hence, we investigated the strengths and weakness of various mechanisms for local and remote distribution. We used the web, podcast distribution as well as distribution services such as Google Video, YouTube and iTunes U. Our goal was to provide guidelines to the university.

For local distribution, we required about 37 GB of storage space per semester. Our lectures consumed about 60 TB worth of network data of which about 6 TB was from users within the campus. The specific values for a particular course depends on its popularity and will likely remain the same regardless of whether the university captured the lecture videos or whether the faculty members personally captured them. However, the ease of our capture can make the cumulative resource requirements more acute. Also, some courses continue to remain popular even after the same instructor offered the course again in a more recent semester. This popularity has significant implications on video archival policies.

For remote distribution, we used the HD video streaming capabilities of YouTube. However, the capabilities offered by free services are dictated by the service provider with little input from the faculty. For example, YouTube does not allow the students to download the videos. The quality and resolution of the free Apple provided iTunes U videos were inadequate for our purposes. Recently, Google Video had also disallowed new video uploads.

Next in Section 2, we describe our personal video capture in further detail. We describe our experiences in video distribution in Section 3 with subjective evaluation in Section 4. We conclude in Section 5.

2 PERSONAL LECTURE VIDEO CAPTURE

2.1 HD to Capture Multiple Modalities

We needed to capture and compose all the class interactions in order to create a single video. Videographers use a variety of means to compose videos. The Camtasia tutorial (Fig. 1a) shows the instructor overlayed with an outline and the slide screencast in a nonoverlapping fashion. On the other hand, the video (Fig. 1b) captured by Rowe and Casalaina [22] overlaid the presenter onto the slides. Friedland and Rojas [23] developed a mechanism to segment the instructor allowing one to carefully overlay them on the slides and further improve the video usability. Customized overlays can reduce the spatial dimensions of the final video while still including all the relevant information. However, creating complex overlays is laborious and hence unsuited for our purpose.

Our primary observation is that a single HD stream captures most of the events in sufficient resolution to be useful for instructional purposes. We require good resolution in order to read blackboard illustrations. Consider a HD screenshot from one of our lectures (Fig. 2) which captures the entire blackboard and the projection while still leaving small notations on the board readable.

To further show the benefits of the HD stream, we overlay rectangles of sizes 720×480 (NTSC DVD) and 320×200 (iTunes U). The insets show the area of capture for
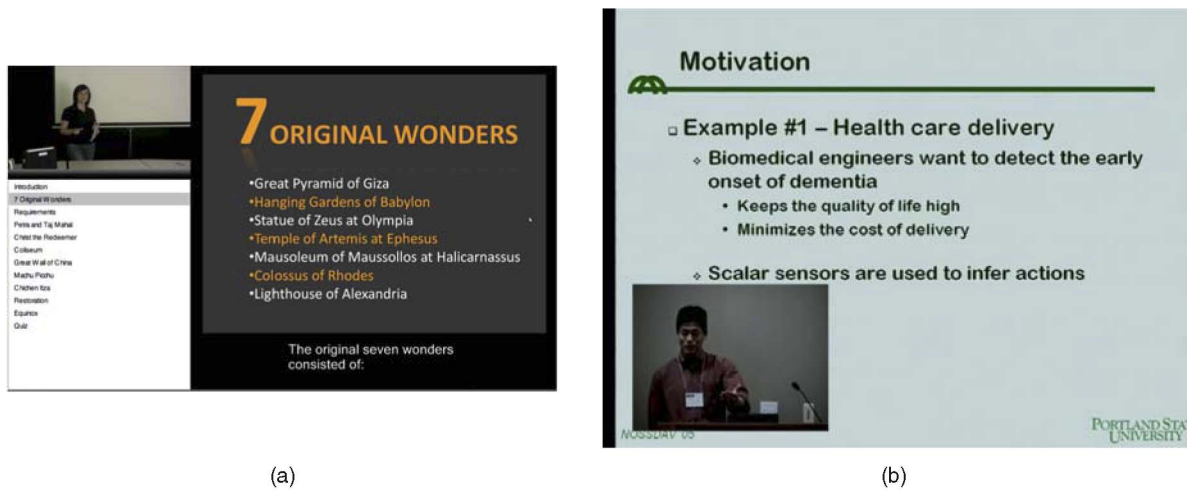
Fig. 1. Composing streams from different modalities. (a) Nonoverlapping. (b) Overlapping.

the various low-quality capture schemes while retaining the HD resolution. A 320×200 stream only captures a narrow region around the instructor's hand that is writing on the blackboard. At this resolution, we require a videographer to manually pan the camera in order to show the relevant details, especially when the instructor continues writing and moves to a different part of the blackboard. The results are slightly better for the 720×480 stream as it covers a larger portion of the image. Commercial motion tracking camera systems are not sophisticated enough to sometimes stop following the instructor and focus on items that the instructor is pointing toward. Recently, Nagai [24] also captured HD video streams of lectures. However, they then created SD videos from these HD streams by automatically tracking the instructor; effectively mimicking the behavior of object tracking cameras. The authors claim that images from a fixed camera do not provide sufficient visual interest

to the viewer. On the other hand, instructors are likely to use sentences such as "Over on the far right corner" to point to important concepts; one requires an attentive video technician who can listen to the lecture and pan to the correct part of the board.

It is also possible to capture this entire scene as a (say) 320×200 stream. We overlaid such an image on the top-right corner of Fig. 2. A 320×200 resolution makes the scene unreadable for instructional purposes; the resolution is sufficient to note that the instructor is writing something on the blackboard but is insufficient to decipher what was actually written on the blackboard. However, for the HD stream, a static camera can capture the entire blackboard, LCD projection, and the instructor without requiring additional effort.

Also, our HD lecture videos exhibit good temporal redundancy (e.g., top of the blackboard). Hence, they
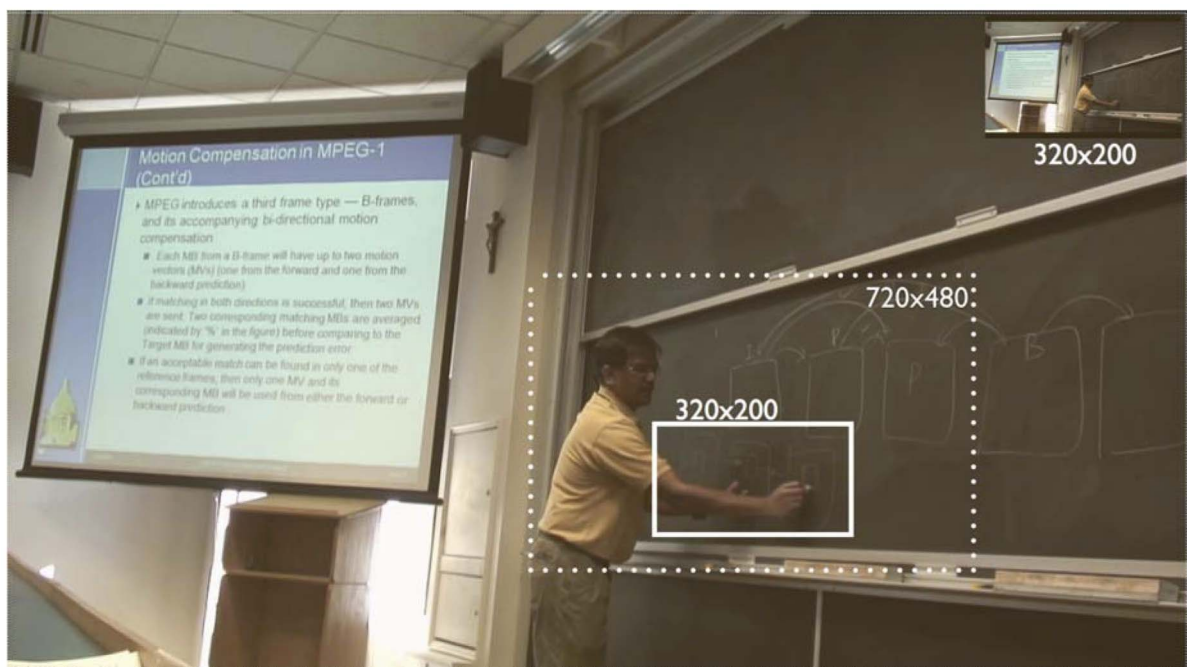

Fig. 2. HD $(1,920 \times 1,080)$ screen resolution. NTSC-DV (720×480) and iTunes U (320×200) illustrated in inset box.
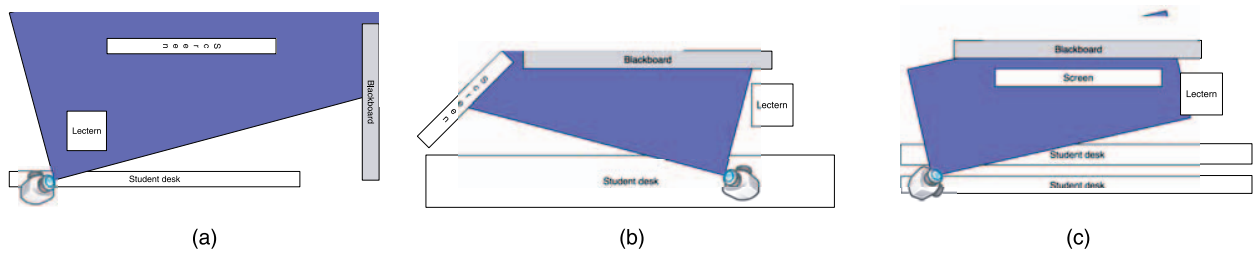
Fig. 3. Lecture hall layouts. (a) Small lecture hall. (b) Medium lecture hall. (c) Large lecture hall.

achieve good compression ratios even when they are spatially larger than videos captured by videographers.

2.2 System Constraints

Our video capture was guided by two design principles:

1. Minimizing the amount of faculty time required for the capture workflow. Some of the times are dictated by the registrar (e.g., time between lectures restricts the time available to set up and pack the equipment), some depend on the technology limitations (e.g., time to transfer video from camcorder) while others are under the control of the instructor (e.g., amount of annotations).
2. Only using commodity, off-the-shelf components. This allows us to minimize costs and maximize the number of faculty who can capture their lectures.

2.3 Our Capture Approach

2.3.1 Capture Equipment

The lecture halls were equipped with a LCD and document projector and a lectern computer; video cameras were not already deployed. Hence, the video capture equipment should be easy to carry, set up, and pack at the lecture hall. Ultimately, what can be captured depends on the weight of the equipment as well as the time required to set up them. This is particularly important because the university allotted duration between lectures is small. Often, prior lectures overran the allotted time; further reducing the time available for set up.

We leverage the low cost advantage of commodity components. We used the Sony HDR-HC1 HDV camcorder (US$1,350 in January 2006) which can record 64 minutes of 1080i HD video (rectangular $1,440 \times 1,080$ pixels) on mini-DV tapes. Depending on the class, our lectures either lasted for 50 or 75 minutes. The Sony HDR-HC1 was one of the first consumer grade HDV cameras. Newer tapeless mechanisms such as AVCHD can store the video in hard disks and flash memory. For example, the Sony HDR-XR150 HD Handycam retails for under US$700 and offers 120 GB of hard drive-based storage that can store up to 50 hours of HD video. These newer camcorders provide adequate storage for lecture capture.

2.3.2 Video Capture Setup

Typically, most of the seats in the classroom were occupied. Finding a location to set up the video camera that provides a good view for video capture while also not obstructing any students from viewing the lecture was challenging, especially since it was impractical to carry tall tripods to

each lecture. The layout of each classroom was different; we used four different types of halls over the past seven semesters; Fig. 3 illustrates the layout of three such lecture halls. The small classroom was flat, the seats in the medium room was elevated while the seats in the large room was steeply elevated to accommodate about 120 students. Flat classrooms require the camera to be installed in the front (unobstructed view) while elevated classrooms require placement further back. We mounted the camcorder on a Manfrotto 209 Tabletop Tripod with a 482 Micro Ballhead (portable, about "4" in height and retails for US$55). The height of this setup was unobtrusive. In general, placing the tripod further back increases the camera field of view. Wider field of view can capture more aspects of the lecture. However, it will also capture (the backs of) some of the students, which is undesirable for our purposes. Note that cameras that are installed by the university and mounted on the ceiling (like in [24]) could be installed further back and still avoid capturing the students.

At the beginning of each semester, we surveyed the lecture hall and chose a good location to place the video camera. Depending on the topics planned for a particular lecture, we tweaked this location. For example, when we expected to use the blackboards much more than the LCD projection, we adjusted the camera to place more importance on the blackboard. In general, the quality of the video was robust against the location choice. Typically, we chose a location toward the end of the first row (Figs. 3a and 3b); we chose the third row in the larger classroom (Fig. 3c). Placing the video camera among students caused the camera to capture student murmurs. We used a bluetooth wireless microphone that directly connected to the camcorder (Sony ECM-HW1) to prevent the camera from capturing student conversations. After the initial choice of a location, we experienced little problems in reusing the same location to place our video camera (students also typically sat in the same location throughout the semester). We made sure that we did not capture any students in the video in order to protect their privacy. We manually removed scenes where students walked into the camera field of view to (say) turn in their assignments.

2.3.3 Capture Experience

Between the Spring 2006 and Spring 2009 semesters, we recorded the lectures of seven courses in four different types of lecture halls. In the Spring of 2006, 2007, and 2008, we recorded the lectures of a junior level Operating Systems course. These classes convened for 50 minutes each, three days a week for a total of about 36 lectures. This was a core required course for all Computer Science majors. This

TABLE 1
Usage Statistics for the Various Lectures (February 2006-November 2009)

Course	Size (GB)	Download amount			Download count		
		All (GB)	Internet (%)	Campus (%)	All	Internet (%)	Campus (%)
All	257.9	60,337.5	90.34	9.66	199,837	94.69	5.31
Spr '06	23.4	18,115.3	97.27	2.73	86,660	96.51	3.49
Fall '06	21.6	3,610.07	87.82	12.18	9,833	90.76	9.24
Spr '07	51.1	15,207.1	93.08	6.92	48,572	96.13	3.87
Fall '07	23.1	3,818.57	93.28	6.72	9,031	93.93	6.80
Spr '08	59.1	13,747.5	79.45	20.55	32,407	88.59	11.41
Fall '08	54.4	5,391.47	86.18	13.82	12,413	95.79	4.21
Spr '09	21.9	447.44	96.64	3.36	921	97.83	2.17

course provided the necessary background for the graduate Operating Systems course which we taught in the Fall 2008 semester. The graduate course was considered to be a qualifying exam and was required of all incoming graduate students. Note that many of the graduate students did not graduate from the university itself; they likely took the undergraduate Operating System course from their own institutions. Some of the graduate students did not hold an undergraduate degree in Computer Science and hence never took an undergraduate Operating Systems course. Regardless, all the graduate students were strongly encouraged to review the course materials covered in our undergraduate Operating System course (especially since graduate students did not receive any graduate level credit for taking the Junior level course). This graduate course met twice a week for 75 minutes each for a total of 26 lectures; the camcorder could only capture about 64 minutes of each of these lectures. In the Fall 2006 and Spring 2009 semesters, we taught an undergraduate Multimedia Systems course which was also cross listed as a graduate course. The Fall 2006 course was offered twice a week for 75 minutes each while the Spring 2009 course was offered thrice a week with each lecture lasting 50 minutes. In the Fall 2007 semester, we also taught a undergraduate/graduate course on Networked Sensor systems. This course met twice a week for 75 minutes per lecture. Note that the videos from classes that met twice in a week were smaller (Table 1) because we only captured 64 minutes of the 75 minute lecture. Newer video cameras that use the AVCHD format will not experience this capture limitation. Also, earlier courses used lower bit-rate high definition videos than what was used in later semesters.

Our primary focus during the lecture was in interacting with the students and not to face and talk into the camera. We only acknowledged the existence of the camera when discussing private information (such as student grades). Sometimes this meant that the lecturer would walk away from the camera or continue writing past the camera's field of view; these events were rare because of the wide capture angle of the camera. Note that a trained technician would have followed our movements and generally did a better capture job. We also did not use any special lighting facilities; the lighting in typical classrooms was adequate for video capture.

During each lecture, we projected Powerpoint slides using the lectern PC. We experimented with the presentation capture feature of Powerpoint. During the postprocessing stages, we can then combine the video and slide capture streams using tools such as Camtasia. Unfortunately, Powerpoint missed the synchronization timing between the audio streams (captured by the camera) and slide transitions (recorded by Powerpoint). It also lost the audio segment if we went back to a previous slide. The university managed lectern PC's did not support the Camtasia tools. We believe that carrying our own laptop with Camtasia tools places an undue overhead in terms of carrying, setting up, and dismantling two devices (camera and a laptop) for each lecture. Also, the university assigns the time between lectures and is limited. In general, it took us about five minutes each to set up and pack-up the video gear (there was usually 15 minute breaks between lectures).

2.4 Postprocessing

The next step is to transfer the video from the camera, perform any editing and annotation operations and convert them into a form that is suitable for distribution. There are a wide variety of video processing options[3] available for Windows, Mac, and Linux operating systems. The platform choice affects our ability to create annotations that are targeted to specific clients. For example, tools from Apple are better integrated to produce annotations for iPod users while YouTube annotations are platform independent. We used Apple products for our postprocessing; these tools were available for free with the purchase of new Apple hardware and require no additional configuration. Regardless of the operating system used, multimedia processing is CPU intensive; the instructor should invest in the fastest possible processor in order to reduce the video processing duration.

After each lecture, we uploaded the videos using the IEEE 1,394 Firewire interface. Throughout the seven semester capture interval, we used different computers to leverage any advances in the capabilities of commodity computers. Initially [25], we used a single core Apple iMac desktop with a G5 2 GHz processor. From Fall 2006, we used an Apple Macbook Pro with a 2.16 GHz Intel Core Duo processor. Recently, we used an Apple Macbook Pro with a 2.26 GHz Intel Core2 Duo processor.

Processor capabilities affects all aspects of the workflow. Consider the time to transfer the video from the

3. http://en.wikipedia.org/wiki/List_of_video_editing_software.

camcorder; the iMac desktop took two hours to download the 50 minute lecture. The other setups performed this operation in real time. The mini DV tapes restricted the download duration to real time. Newer AVCHD camcorders allow faster transfer using the 480 Mbps USB cable. Initially we used the Apple Finalcut tool for video editing. However, we almost immediately switched to the Apple iMovie software for its simplicity. The iMovie software is also distributed freely with a new Mac. The videos can also be processed by Camtasia which offers a richer set of composition options (requiring more instructor time). Once transferred from the camera, a 50 minute lecture requires over 30 GB of raw storage. Though this size was not a concern for storing videos on a modern laptop (640 GB of laptop disk retails for US$90), this size was impractical for distribution to the students.

Once the video was transferred to the computer, we manually annotated the videos. The specific annotation depends on the amount of time that the instructor was willing to spend on creating them. It also depended on the approach to distribute the video and is discussed in detail further in Section 3. After the annotations were performed, we transcoded the stream into three different formats: 1) a one Mbps HD video (more recently increased the stream bandwidth requirements to two Mbps in order to leverage technology improvements) encoded using H.264 with a screen resolution of $1,280 \times 720$, 2) a video object customized for a video iPod/iPhone—H.264 stream that was initially encoded at a resolution of 320×180 and later at a resolution of 480×270 (to account for improvements in the ability of students to consume higher fidelity videos), and 3) a MPEG-4 audio podcast created using the Apple Garageband tool. Garageband allowed us to create slide markers, attach Powerpoint slide images to the slide markers and add text annotations. Note that the iPod video can be played using the Quicktime player on a computer as well as the Sony PSP game gear. Initially, transcoding to the iPod video formats took around 3-4 hours while the HD video took about 10-12 hours. More recently, these operations take less than a quarter of these times.

2.5 Summary of Lecture Video Capture

Lecturers were typically skeptical of the time and effort required to capture every lecture themselves. Our experience shows that these costs are modest. The HD camcorder and accessories cost us under a thousand dollars and are likely usable for a few years. Before the beginning of the semester, the faculty are required to survey the classroom in order to choose the camera placement; the location choice is robust. The fixed cost for capturing lectures include the 1) time to set up and pack-up the cameras during each lecture, 2) time to transfer the video from the camera, and 3) time to transform the videos into formats that are distributable to the students. The setup times are limited by the durations between lectures. The time to transfer the videos is becoming faster while the time to transform the video benefits from continuous improvement in processor capacity. Both these operations do not require active involvement of the instructor. Instructors can control the amount of time spent on annotating the videos; elaborate annotations can take a long time.

3 DISTRIBUTION OF LECTURE VIDEOS

The next step is to choose the distribution mechanism. In our university, each course is alloted one GB of storage for instructional needs. Unlike review materials such as Powerpoint slides, videos are large and require significant amounts of storage. On average, we required about 37 GB of storage per semester (Section 3.1.2); the current storage allocation is inadequate. Even though the high storage requirements will remain the same regardless of who captured the videos, personal capture forced us to investigate the various storage and distribution options. The goal is to gain insights and ultimately partner with the campus technology support personnel in order to choose the mechanism that is appropriate for the entire campus. We identified three challenges:

1. *Storage cost.* Unlike prior video capture mechanisms that require a videographer to be physically present in each lecture, the storage support personnel need not be in-situ. However, the cost to expand traditionally managed storage to accommodate all the videos is nontrivial. Each semester, our 2,500 courses in the entire university would require 92 TB of storage; providing a reliable and managed storage for this amount is expensive. Just a few years ago, the university allocated an order of magnitude less storage per course. If those trends continue, the university might ultimately invest in enterprise class storage for storing videos. In the meanwhile, we boot strap the process by arguing for a storage solution that relaxes traditional reliability guarantees. We expand on this storage in Section 3.1.1.

2. *Local versus remote distribution.* The student location plays an important role. Current students access the lectures from campus, dormitory as well as from off-campus locations. Currently, the university uses a 200 Mbps link to access the Internet. The university also has special peering agreements with some local ISPs in order to service students who live off-campus. Also, private email conversations show that our alumni are continuing to access the videos from remote locations. Hence, we investigate local as well as remote distribution.

3. *Public versus private.* Another question is whether the videos should be publicly available or restricted only to the students who took the course. Maintaining access control lists, especially for alumni can be hard. Hence, we publicly released the videos to everyone. Publicly releasing the videos meant that the number of accesses could be high. For example, Camtasia tools can use the Screencast[4] service to distribute videos. Screencast provides 25 GB of storage and 200 GB of transfer bandwidth for about US$9.95 a month with an additional US$31.95 per 100 GB transfer block. Each of our HD lectures consumed about 1.25 GB of storage (processed using Camtasia). The access costs can quickly accumulate.

We distributed the videos using YouTube, Google Video, and a local web server. Next, we describe how the videos

---

4. http://www.techsmith.com/screencast.asp.

were accessed using these approaches. Also, storage for older lecture videos might need to be reclaimed; we analyze the long term popularity of lecture videos.

3.1 Distributing Videos from Inside the University

We distributed the videos from our own web server. A variety of web server software is readily available for all the major operating systems. The videos can be downloaded from the course web page or using podcast feeds. Rather than choosing enterprise level servers, we chose an entry level server. Using two 7,200 RPM 2 TB hard disks (retails for US$120 each) in a mirroring configuration offers enough space to store the videos of over 50 classes. It is cheaper to maintain redundant copies on hard disks rather than use backup tapes. The video objects follow a write-once, read-many access model. Users also download the videos rather than stream them. Hence, the large read throughput dominates our workload; contemporary entry level servers offer sufficient capacity. We envision scaling this setup to the entire university by roughly provisioning one server for each department. Older lectures can be archived by configuring a new server at the beginning of each year. Universities can gradually replace each new department level server with more enterprise level hardware.

We distributed the videos using the MPEG4 [26] format. MPEG4 players are widely available for desktop as well as mobile users (e.g., Apple iPod/iPhone, Microsoft Zune, and Sony PSP handheld game units).

The next challenge is to publicize the location of these videos. We published the URL of the lecture videos on the course web page. Additionally, we created a podcast of the lecture videos using the freely available Vodcaster tool.[5] The podcast is a web syndication mechanism that uses a Really Simple Syndication (RSS) XML feed to point to the web location of the videos. Students subscribe to the feed using programs such as Apple iTunes, Microsoft Zune or directly from their Apple iPhone/iPad. These programs frequently query the feed in order to discover new videos. The client programs can be configured to automatically download the latest videos for offline use. An Arbitron study [27] showed that over 23 million Americans used podcasts in January 2008. Previous work [1], [28], [29], [30] had also discussed the ease of using podcasts for instructional purposes.

In the Fall 2008 semester, we also distributed the lectures using Apple iTunes U.[6] iTunes U appears similar to podcasts; students subscribe to the lectures from the university's iTunes U section. However, Apple allows the instructors to customize the way that the page appears to the user. For example, instructors are allowed more control over the related links. Instructors can also organize the objects using tabs. For our lectures, we distribute the lecture slides in PDF format as well as audio and video format of the lectures. Podcasts display them on a *most recent first* basis. However, iTunes U allows the instructor to organize each object in its own tab. We illustrate our iTunes U page in Fig. 4b. Unlike the screen for the corresponding podcast (Fig. 4a), the iTunes U allows us to organize the various

objects as *Slides*, *Audio*, *Video*, and *Assignments*. On the other hand, the video objects themselves are served from Apple servers. Apple provides the storage and distribution resources for free to educational institutions. However, at the time of this writing, Apple automatically downgrades the videos to a 320×180 resolution low quality video which loses many of the benefits of our high definition videos.

3.1.1 Video Annotations Useful for Local Distribution

An important feature of personal capture is the ability of the instructor to add meaningful annotations post hoc while processing the videos. The instructor can splice the video and add a new video clarification. They can add markers for slide transitions as well as overlay textual clarifications. Video editing software make it relatively easy to add these annotations. Annotations which modify the video are available to the student regardless of the distribution mechanism. However, certain annotations depend on the distribution mechanism. Note that the time required to add annotations directly depends on its complexity; the instructor should strike the right balance.

For local distribution, we used annotations that are viewable on the iPod as well as on the Quicktime player. For the video objects, we manually marked the time at which we changed the Powerpoint slide (on the LCD projection). For the audio objects, we added a still image that showed the Powerpoint slide that was being discussed. These annotations appear differently in different players. For example, the audio podcasts can show the slide markers (Fig. 5b) or the slide images themselves (Fig. 5a). Playing the audio podcasts via Quicktime shows the slide images and chapter markers (Fig. 5f). On the other hand, the video podcasts can display the slide markers (Fig. 5d) as well as the actual video (Fig. 5c). These annotations allow the students to choose the appropriate component of the video for quick review. Note that these annotations will not be visible if the audio and video objects are viewed on a player which did not recognize them, such as in the Sony PSP handheld unit.

3.1.2 Usage Statistics for Local Distribution

First, we tabulate the amount of data transferred as well as the number of audio and video objects downloaded between February 2006 and November 2009 in Table 1. We also show the percentage of requests from within the campus as well as from the public Internet. As we noted in Section 2.3, the amount of data created in some semesters was smaller because of the 64 minute capture limitation. We serviced about 60 TB worth of data for over 200,000 objects. Of these, about 9.66 percent of the data (8.8 TB) were requested by on-campus users while the remaining 54.5 TB of data were requested by Internet users. Assuming a network capacity of 200 Mbps to the Internet, external users consumed videos worth over 25 days of our external network connection.

Analyzing the data for objects created for the different classes, we note that some classes were more popular than others. For example, the Spring 2006 offering serviced over 86 thousand requests (as compared to 200 thousand requests for all the semesters). In general, all the undergraduate Operating System courses were popular and serviced over 167 thousand requests (84 percent of all

(a)

(b)

Fig. 4. Screen capture of podcast and iTunes U. (a) Podcast (shows all the available objects in a single screen). (b) iTunes U (customized tabs for *Slides*, *Audio*, *Video*, and *Assignments*).

requests) and used about 47 TB (78 percent of the transferred data). Among campus users, the graduate Operating System course (Spring 2008) was popular, accounting for 20.55 percent of the data for that course. The Multimedia system course offering was also popular.

Next, we plot the quarterly change in the popularity of the various classes, both from inside the campus and from Internet users in Fig. 6. We observe a flash crowd in the second quarter of 2007 for the Spring 2007 Operating Systems course. Earlier, Table 1 showed that Spring 2006 course was popular. Among Internet users, Fig. 6 shows the popularity of the Spring 2006 offering increasing from serving 1.8 thousand objects in the second quarter of 2006 to over 9.4 thousands objects by the second quarter of 2009. Even among the campus users, the popularity remained stable at around 0.2 thousands. Note that the campus users exhibit a seasonal variation between summer and the rest of the academic year; the school does not offer many courses over the summer break. We observe that most lectures

continue to remain popular, especially since the recent course offering in Spring 2008 could potentially subsume similar courses offered in the Spring of 2006 and 2007. It is likely that students who took the Spring 2006 offering preferred to review using those videos instead of using videos from the newer offerings of the same course. We observed that lectures are a continuum, replacing a single lecture from one semester with the corresponding lecture from a prior offering is not straightforward. One mechanism to conserve resources is to stop servicing requests for older courses. If users continue to request older offerings, we believe that objects should not be expired—at least within the three year window used in our analysis.

Finally, we illustrate the quarterly change in resource consumption for audio, SD, and HD videos in Fig. 7; Fig. 7a shows the magnitude of change both as a count as well as the amount of data transferred while Fig. 7b shows the relative percentage of each type of object. From Fig. 7b, we note that the relative popularity of audio objects is waning, in terms of
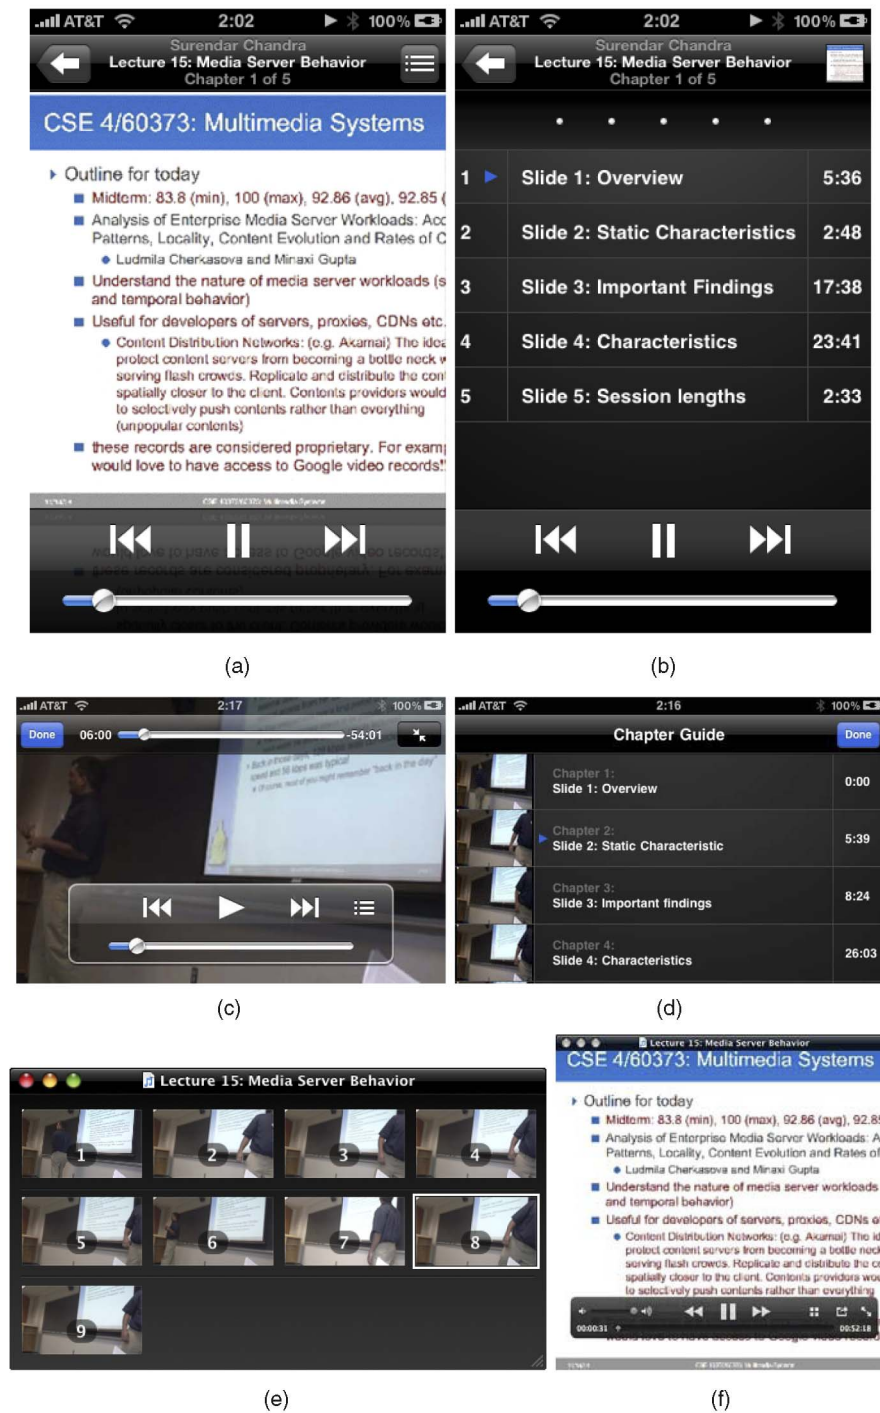
Fig. 5. Enhanced audio/video podcasts (with slide markers and slide images. (a) Annotated audio on an iPhone. (b) Chapter markers for audio. (c) Annotated video on an iPhone. (d) Chapter markers for video. (e) Desktop quicktime player—video. (f) Quicktime player—audio.

volume: from about 10 percent in the second quarter of 2006 to 3 percent in the third quarter of 2009 and in terms of count: from about 25.8 to 22 percent, respectively. Grabe and Christopherson [31] also observed that psychology students did not prefer audio. The SD videos became inexplicably popular in the second quarter of 2007. Though such flash crowds are common in Internet scenarios, the size of the audio and video objects place tremendous stress on our networking infrastructure. Interestingly, HD videos are becoming more popular; having increased in count from 13.2 to 29.6 percent with the corresponding data volume

from 21 to 59 percent. One of the persistent student complaints in Spring 2006 was the enormous size of HD videos; commodity technologies appear to be evolving to allow more students to use the HD videos. We saw corresponding drop in the popularity of SD videos. However, there is little evidence that our campus Internet connection is scaling at a similar rate to accommodate the three fold increase in the volume of HD videos.

In terms of the absolute counts and the amount of data transferred (Fig. 7a), we note a steady increase in the amount of data transferred in each quarter. The amount of
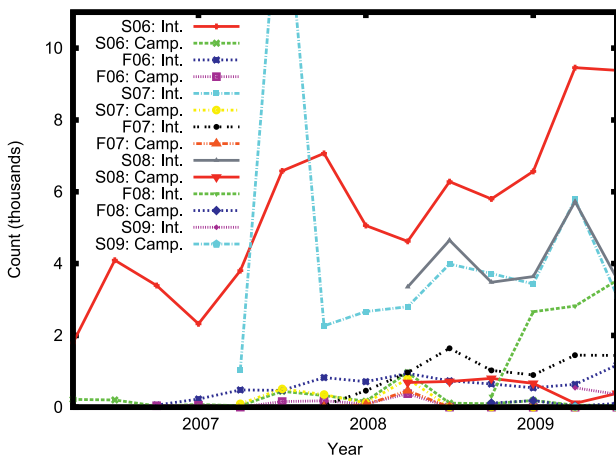
Fig. 6. Popularity of lectures by the semester.

data consumed in a quarter by the HD videos increased from 0.3 TB in 2006 to over 5.6 TB. During the flash crowds in 2007, the SD videos also consumed about 5 TB of data in a single quarter. By 2009, we were consuming 9.4 TB in a single quarter or around 4.4 days worth of campus Internet connectivity.

Even though the University does not currently limit the amount of network resources used by a faculty member, the level of resource usage highlighted in this section is not sustainable, especially when other faculty members also release their videos for public consumption. The author recently participated in the university iTunes U advisory panel. Apple allows the university to store 500 GB worth of data on its cloud servers. The university can also host videos on its own servers. Many faculty and administrators of the panel assumed that the primary difficulty in having an iTunes U presence for the university is in producing the content for distribution over iTunes U. Unless the individual faculty member objected, there was unanimous support for publicly releasing as much contents as possible. However, our experience suggests that the cost of personally creating the video contents was relatively small. However, the storage and distribution costs can quickly overwhelm the campus resources if a significant fraction of the faculty followed in the author's foot steps and personally captured and distributed their own lecture videos.

3.2 Distributing Videos from Outside the University

Given the cost to the university for distributing HD videos to students who reside outside the campus network, we investigated distribution mechanisms that stored the videos outside the campus. There are two classes of paid distribution mechanisms: streaming services such as Screencast charge a monthly fee for the storage as well for the network bandwidth used for streaming the videos. Cloud services such as Amazon S3[7] also offer a viable alternative for storing and distributing objects. Paid services allow the instructor to service the videos without advertisement banners. However, the 3 TB of network resources used recently by our lecture videos (Fig. 7a) will cost about $550/month at Amazon. Hence, we investigate free (i.e., advertisement supported) services.

7. https://s3.amazonaws.com/.

There has been a proliferation of free video hosting services. However, many of these services only allow videos of short durations. Google Video became available right during the Spring 2006 course. Google Video did not restrict the length of the video segment. Hence, we used Google to distribute the lectures captured for the six semesters between Spring 2006 and Fall 2008. However, Google recently discontinued video uploads both for free users as well as for Google education premium users. Hence, we investigated YouTube. Initially, YouTube restricted videos to 100 MB. However, YouTube allowed longer uploads for *Director* level members. During the Fall 2008 semester, they increased the upload limitation to 1 GB and have since raised them to two GB. Since November 2008, YouTube supported HD streaming at $1,280 \times 720$ resolution. On November 12, 2009, YouTube announced support for $1,920 \times 1,080$ resolution. With this addition, YouTube is a useful platform for our purposes. We made the Fall 2008 and Spring 2009 semester contents available in the HD format. We also made the Spring 2008 videos available in SD video format. Recently, YouTube has discontinued new enrollments to the *Director* program; educators can upload longer videos through the YouTube EDU program. As a free service, one is limited by the vagaries of policies set by the video distributors.

3.2.1 Video Annotations Useful for Remote Distribution

We describe our experiences with streaming as well as annotating videos using the YouTube service (Google Video did not support annotations). Note that we do not have control over the annotation mechanism or the policies on whether the object can be downloaded. For example, YouTube does not allow the students to download the videos; students are expected to be online while watching the stream. Given the proliferation of smart phone and laptops that are capable of playing YouTube streams, this restriction might be acceptable.

As a free service, the specific annotation mechanisms are controlled by YouTube and are evolving continuously. The annotations are browser based and are available from a wide variety of browsers and operating systems. YouTube allows a rich set of annotation that uses *Speech bubble*, *Note*, and *Spotlight* to directly add annotation elements into the stream at a specified time and spatial location. The instructor can also control the font and color elements in these annotations. The instructor can also authorize other users to annotate the videos. However, the system does not report the provenance records on where any annotations were made. Hence, we did not use this feature for our lectures. Even though these annotations are powerful, we believe that they are inadequate for instructional purposes. It is not possible to index and list all the annotation elements in a video, the annotations are viewed when the user watches the particular video segment. Lecture videos are not always watched sequentially; students require the ability to jump to discussions about specific slides, a capability already available from our local distribution (Section 3.1.1). Regardless, we continue to explore ways in which we can utilize annotations on YouTube.

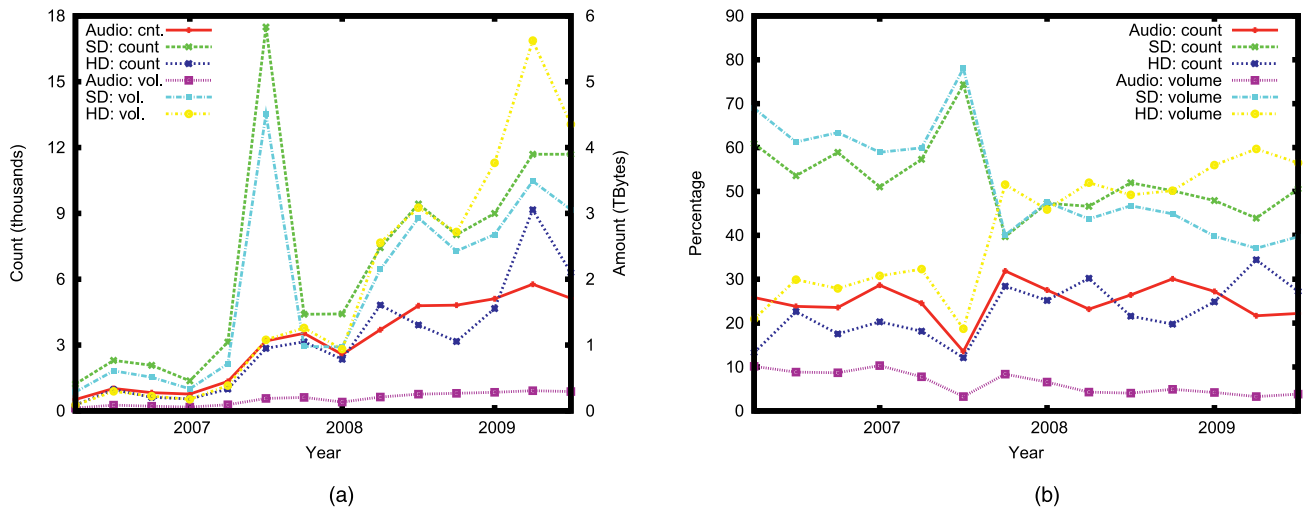(a)   (b)

Fig. 7. Change in popularity of the audio, SD, and HD videos. (a) Magnitude. (b) Percentage.

3.2.2 Usage Statistics

We plot the number of accesses as well as their geographical origin (as reported by YouTube) in Fig. 8. From Fig. 8a, we note that the number of accesses are increasing with over 200 access per day by February 2010. Also, in Fig. 8b, the darkness of the state indicates the popularity of the requests from that state. Most requests came from Indiana, the location of the University. A large number of requests also came from Ohio, a neighboring state as well as from California. California is a popular job destination for Computer Science graduates. It is possible that most of the requests from Indiana are from inside the campus, which defeats the purpose of making the videos available to

(a)

(b)

Fig. 8. Usage statistics from YouTube. (a) Number of unique users and total views. (b) Geolocation of viewers from the US.

Internet users. On the other hand, serving users from Ohio and California from YouTube can reduce the network load on the campus Internet link. Incidentally, these requests from YouTube have not made a significant impact on the number of requests from the campus (Fig. 6).

3.3 Summary of Distribution Related Issues

We showed the vast amounts of network resources required to service the video objects as well as their enduring popularity. Recent improvements in the quality of videos serviced by YouTube allows the instructor to distribute HD videos for free; important, especially when the university was not providing the required storage and distribution infrastructure. Ultimately, universities can use our experience to strike a balance between local and remote distribution and trade off distribution cost with the control afforded by local hosting. Distributed storage solutions also allow the university to incrementally scale up the storage volume.

4 SUBJECTIVE EXPERIENCE

Several projects at different universities have distributed lecture videos. Next, we describe our own experiences with capturing the lecture videos, both from the perspective of other faculty as well as through student feedback.

4.1 Faculty Concerns

The primary faculty concern was that students would not attend class. Prior reports on this count had been mixed. Rowe et al. [9] note that 30 percent of Berkeley students did not attend lectures whether the class was webcast or not. Harpp et al. [4] observed a small drop (about 10 percent) in student attendance for screencasting their lectures. Similarly, Copley [32] observed minimal drop in student attendance. Traphagan et al. [33] observed a drop even though the drop was steeper for distributing the Power-point slides. They also observed improvements in student learning experience. However, we did not observe any drop in student attendance. Student feedback offers an explanation for this behavior. Our students have a busy schedule. Skipping the class meant that they needed to find another time to listen to the videos. Unless there were some
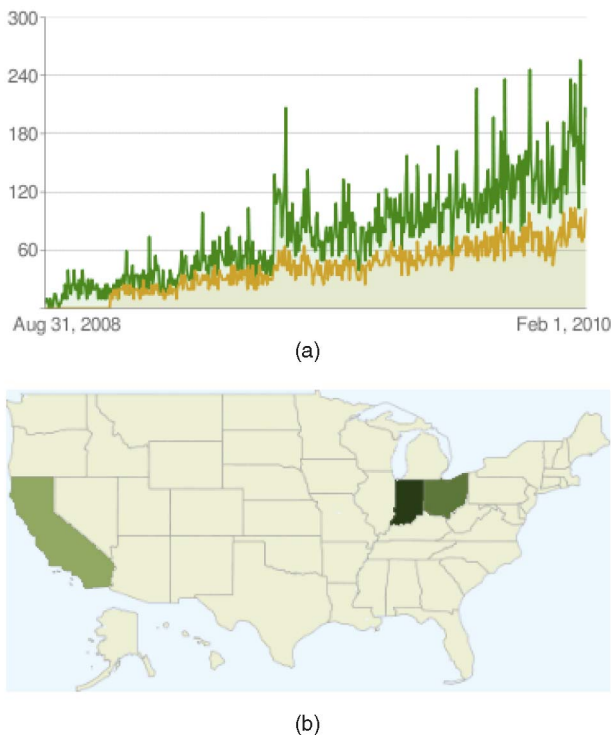
extenuating circumstances, it was better to attend the lecture. Watching a stored video reduces the penalty for not attending a lecture but does not reduce the cost of actually listening to a lecture.

Also, video recording leaves a record of every misspoken or incorrect words uttered by the faculty. Students can use them to confront the faculty (difference between *I think you said that "1 == 2"* versus *You said "1 == 2" on 24 Feb. 2006 at 10:54:23 AM*). Personally, we consider this to be an acceptable risk. Faculty are not infallible; they do not have to act otherwise. However, they might discard unsubstantiated criticism from anonymous YouTube users.

The other faculty concern was that this will take up too much precious time without any tangible benefit to the students. Our analysis shows that the videos remain popular for over three years even among the local campus users; the effort is worth the hardship.

The final faculty concern was about the intellectual property implications of such recordings. Clearly the university holds the rights to all the lectures. Some schools restrict the distribution of distance learning videos to students who had registered for the course. Our university does not offer such courses and so has no explicit policy that governs video dissemination. The recent efforts by the university to produce contents for iTunes U suggest that the university was willing to distribute the videos for free. On the other hand, the laws concerning video distribution of material that were shown in the classroom under the *fair use* doctrine is myriad. The instructor should consult with the university counsel regarding their legal obligations.

4.2 Student Feedback

The student feedback had been positive with no observable drop in student attendance. Several students expressed the view that they preferred the organized class setting over a chaotic dorm. However, one student who suffered from anxiety disorder found it more convenient to entirely watch the videos. Of course, a video was the only option when the instructor or the student was traveling. One student mentioned that when he dozed off in class and woke up, he made it a point to note down the exact time that he woke up so that he can go back to the materials that he missed.

Students reported archiving written lecture notes (the author has a pile of decade old notes). However, with the passage of time, these printed notes loose their context. Several of our students archived the lecture videos in a DVD along with the printed notes. Several students wished that they had the videos from their own *Linear Algebra* courses. They noted that it is not helpful to sit in on another *Linear Algebra* course taught by a different instructor because they were looking to refresh some specific content that they learned, which may not be taught exactly the same way by every instructor. These observations motivated our effort.

Some alumni who had graduated and joined the workforce reported that they recently watched the lecture videos. They were able to better understand the lectures (e.g., video compression algorithms) in the context of their current work than when they were students at the university.

4.3 Summary of Subjective Experience

Our experience showed that faculty captured videos are as effective as videos captured by videographers or by using automated mechanisms. Students reported their appreciation for the availability of the lecture videos in the course review forms. They described various ways in which they found these videos useful, both while they were a student and even after they had graduated.

5 CONCLUSIONS

There has been considerable evidence on the importance of lecture review videos. Faculty members prefer a fully automated video capture and distribution mechanism. However, automatic video capture mechanisms are not always available in an easily deployable form. Many universities are unwilling to bear the cost and deploy video capture options for all their lectures. Instructors who are convinced of the usefulness of videos must still depend on the university to allocate its scarce resources in capturing their own lectures. Instead, we show that technology improvements allow any instructor to capture and produce the videos with minimal effort. The technology trends are also allowing students to consume HD videos. We showed that the next challenge was in choosing the distribution mechanism which balances the desires of the instructor to freely distribute the video and the strain that their choice can place on the campus network. We offer our experiences that can allow the campus IT personnel to customize a distribution solution that is suited to the location of their students.

REFERENCES

[1] G. Campbell, "There's Something in the Air: Podcasting in Education," *EDUCAUSE Rev.,* vol. 40, no. 6, pp. 32-47, Nov. 2005.
[2] W. Hürst and W. Waizenegger, "An Overview of Different Approaches for Lecture Casting," *Proc. IADIS Int'l Conf. Mobile Learning,* July 2006.
[3] D. McKinney, J.L. Dyck, and E.S. Luber, "ITunes University and the Classroom: Can Podcasts Replace Professors?" *Computers and Education,* vol. 52, no. 3, pp. 617-623, Apr. 2009.
[4] D.N. Harpp, A.E. Fenster, J.A. Schwarcz, E. Zorychta, N. Goodyer, W. Hsiao, and J. Parente, "Lecture Retrieval via the Web: Better than Being There?" *J. Chemical Education,* vol. 81, no. 5, pp. 688-690, May 2004.
[5] J. Lanir, K.S. Booth, and L. Findlater, "Observing Presenters' Use of Visual Aids to Inform the Design of Classroom Presentation Software," *Proc. 26th Ann. SIGCHI Conf. Human Factors in Computing Systems (CHI '08),* pp. 695-704, Apr. 2008.
[6] J.A. Brotherton and G.D. Abowd, "Lessons Learned from eClass: Assessing Automated Capture and Access in the Classroom," *ACM Trans. Computer-Human Interaction,* vol. 11, no. 2, pp. 121-155, 2004.
[7] G.D. Abowd, "Classroom 2000: An Experiment with the Instrumentation of a Living Educational Environment," *IBM Systems J.* vol. 38, no. 4, pp. 508-530, Oct. 1999.
[8] S. Mukhopadhyay and B. Smith, "Passive Capture and Structuring of Lectures," *Proc. ACM Multimedia Conf.,* pp. 477-487, 1999.
[9] L.A. Rowe, D. Harley, P. Pletcher, and S. Lawrence, "BIBS: A Lecture Webcasting System," Technical Report 2001-160, Berkeley Multimedia Research Center, June 2001.
[10] E. Machnicki and L.A. Rowe, "Virtual Director: Automating a Webcast," *Proc. Multimedia Computing and Networking (MMCN '02),* vol. 4673, pp. 208-225, Jan. 2002.
[11] Q. Liu, Y. Rui, A. Gupta, and J.J. Cadiz, "Automating Camera Management for Lecture Room Environments," *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI '01),* pp. 442-449, 2001.

[12] Y. Rui, A. Gupta, and J. Grudin, "Videography for Telepresentations," *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI '03)*, pp. 457-464, 2003.

[13] C. Zhang, Y. Rui, J. Crawford, and L.-W. He, "An Automated End-to-End Lecture Capture and Broadcasting System," *ACM Trans. Multimedia Computing Comm. Applications*, vol. 4, no. 1, pp. 1-23, 2008.

[14] P. Ziewer, "Navigational Indices and Full Text Search by Automated Analyses of Screen Recorded Data," *Proc. E-Learn*, Nov. 2004.

[15] D.M. Hilbert, T. Turner, L. Denoue, and K. Sankarapandian, "Autonomous Presentation Capture in Corporate and Educational Settings," *Proc. IADIS e-Learning*, July 2008.

[16] R. Müller and T. Ottmann, "The "Authoring on the Fly" System for Automated Recording and Replay of (Tele) Presentations," *Multimedia Systems*, vol. 8, no. 3, pp. 158-176, 2000.

[17] S. Repp, A. Groß, and C. Meinel, "Browsing within Lecture Videos Based on the Chain Index of Speech Transcription," *IEEE Trans. Learning Technologies*, vol. 1, no. 3, pp. 145-156, July-Sept. 2008.

[18] J. Adcock, M. Cooper, L. Denoue, H. Pirsiavash, and L.A. Rowe, "Talkminer: A Lecture Webcast Search Engine," *Proc. ACM Multimedia*, pp. 241-250, Oct. 2010.

[19] B. Hubbard, J. Stern, and M. Hancock, "Webcast.Berkeley Launches New Automated Lecture Capture System," http://inews.berkeley.edu/articles/Fall2008/webcast, Sept. 2008.

[20] B. Burdet, C. Bontron, and P.-Y. Burgi, "Lecture Capture: What Can Be Automated?" *EDUCAUSE Quarterly*, vol. 30, no. 2, pp. 40-48, 2007.

[21] W. Hürst, "Indexing, Searching, and Skimming of Multimedia Documents Containing Recorded Lectures and Live Presentations," *Proc. ACM Multimedia*, pp. 450-451, 2003.

[22] L.A. Rowe and V. Casalaina, "Capturing Conference Presentations," *IEEE Multimedia*, vol. 13, no. 4, pp. 76-84, Oct. 2006.

[23] G. Friedland and R. Rojas, "Anthropocentric Video Segmentation for Lecture Webcasts," *J. Image Video Processing*, vol. 8, no. 2, pp. 1-10, 2008.

[24] T. Nagai, "Automated Lecture Recording System with Avchd Camcorder and Microserver," *Proc. 37th Ann. ACM SIGUCCS Fall Conf. (SIGUCCS '09)*, pp. 47-54, 2009.

[25] S. Chandra, "Lecture Video Capture for the Masses," *Proc. 12th Ann. SIGCSE Conf. Innovation and Technology in Computer Science Education (ITiCSE '07)*, pp. 276-280, June 2007.

[26] Int'l Organization for Standardization (ISO/IEC), "Overview of the MPEG-4 Standard," http://mpeg.chiariglione.org/standards/mpeg-4/mpeg-4.htm, Mar. 2001.

[27] A.M.R. Study, "Internet and Multimedia 2008 Reports," Apr. 2008.

[28] P. Ractham and X. Zhang, "Podcasting in Academia: A New Knowledge Management Paradigm within Academic Settings," *Proc. ACM SIGMIS Conf. Computer Personel Research: Forty Four Years of Computer Personnel Research: Achievements, Challenges and the Future (SIGMIS CPR '06)*, pp. 314-317, 2006.

[29] D.J. Malan, "Podcasting Computer Science E-1," *Proc. ACM Symp. Computer Science Education*, Mar. 2007.

[30] F.M. Corl, P.T. Johnson, M.R. Rowell, and E.K. Fishman, "Internet-Based Dissemination of Educational Video Presentations: A Primer in Video Podcasting," *AJR Am. J. Roentgenology*, vol. 191, no. 1, pp. W23-W27, July 2008.

[31] M. Grabe and K. Christopherson, "Optional Student Use of Online Lecture Resources: Resource Preferences, Performance and Lecture Attendance," *J. Computer Assisted Learning*, vol. 24, no. 1, pp. 1-10, Feb. 2008.

[32] J. Copley, "Audio and Video Podcasts of Lectures for Campus-based Students: Production and Evaluation of Student Use," *Int'l Innovations in Education and Teaching*, vol. 44, no. 4, pp. 387-399, Nov. 2007.

[33] T. Traphagan, J.V. Kucsera, and K. Kishi, "Impact of Class Lecture Webcasting on Attendance and Learning," *Educational Technology Research and Development*, vol. 58, no. 1, pp. 19-37, June 2009.

Surendar Chandra received the PhD degree in computer science from Duke University. He held positions in academia at the University of Georgia and Notre Dame and in industry at the FX Palo Alto Laboratory. His research interests include experimental systems topics in multimedia, storage, security, networks, and sensor systems. He is the recipient of a US National Science Foundation CAREER Award and is a senior member of the ACM.