

Designing Presentations for On-Demand Viewing

Liwei He, Jonathan Grudin, Anoop Gupta

Microsoft Research

Redmond, WA 98052

{lhe, jgrudin, anoop}@microsoft.com

ABSTRACT

Increasingly often, presentations are given before a live audience, while simultaneously being viewed remotely and recorded for subsequent viewing on-demand over the Web. How should video presentations be designed for web access? How is video accessed and used online? Does optimal design for live and on-demand audiences conflict? We examined detailed behavior patterns of more than 9000 on-demand users of a large corpus of professionally prepared presentations. We find that as many people access these talks on-demand as attend live. Online access patterns differ markedly from live attendance. People watch less overall and skip to different parts of a talk. Speakers designing presentations for viewing on-demand should emphasize key points early in the talk and early within each slide, use slide titles that reveal the talk structure and are meaningful outside the flow of the talk. In some cases the recommendations conflict with optimal design for live audiences. The results also provide guidance in developing tools for on-demand multimedia authoring and use.

Keywords

Video on-demand, streaming media, digital library

INTRODUCTION

Improvements in network bandwidth, computer performance, and compression technologies are facilitating the routine use of video-on-demand in workplace, home, and educational settings. Entertainment and education are two principal uses of video-on-demand. In this paper we focus on videotaped lectures and presentations of the kind used in distance education, internal corporate training, executive briefings, product proposals, marketing analyses, sales pitches, usability study reports, and so forth. Emerging technologies for time-compressing and skimming digitized audio and video could greatly increase the benefits of viewing informational material on-demand rather than live.

The simplest approach is to make available a digitized recording of the audio and video. However, value can be added by post-processing. For example, one can integrate slides, demos, audience questions, video clips, and references to related material or web sites. A table of contents can be constructed that links these materials. However, post-processing introduces a cost.

As audio-video presentations appear on academic and corporate intranets and Internet sites (e.g., <http://stanford-online.stanford.edu/>, <http://murl.microsoft.com>), several questions arise: 1) Will busy people access video materials from their desktops? 2) If value-added features are provided, will they be used? 3) How should presentations be designed for on-demand viewing? 4) Will different design principles apply, or will they reflect traditional principles of presentation design? 5) If differences are found, can designs be optimized to serve both audiences?

As illustration and analogy, consider the case of online viewing of text. Morkes and Nielsen [13, 15] observed that people are much more likely to skim text on the web. When authoring for web viewing, it is crucial to present key messages first, followed by details and background information; to emphasize titles and headers; and to present one idea per paragraph. These are principles of good writing in general, but they become much more important to readers who are skimming quickly.

This paper partially answers these questions. The Microsoft Technical Education Group (MSTE) has supported on-demand viewing for two years and maintained detailed logs of viewer behavior for research purposes. Videos have been accessed on-demand by over 9,000 people, answering the first question. We can learn much more about this new resource by analyzing viewer activity patterns.

The paper is organized as follows. The next section reviews the costs and benefits of on-demand video, and prior work. Next we describe our system and data collection methodology. The following sections present the results of the data analysis, including general patterns of on-demand viewing as well as those specific to our hypotheses. The discussion summarizes design lessons that we distilled. The conclusion outlines further considerations for the design of presentations and tools to support their authoring.

ON-DEMAND VIDEO: COSTS AND BENEFITS

Digitized video enables large audiences to view the content anywhere, anytime, and accompanied with relevant materials. In addition, viewers can selectively time-share or focus on relevant materials by pausing, fast-forwarding, reviewing, or jumping from segment to segment. They can also quit without risk of offending a speaker.

These benefits must be balanced against the cost of making talks available on-demand, which include the production costs of recording, digitizing, and post-processing talks for on-demand delivery, and the cost of video-servers to store

the talks. Indirect costs include increased network usage and possible infrastructure upgrades. Another disadvantage of distributed video, whether real-time or on-demand, is that it generally does not allow direct interaction between the speaker and the distant audience (but see [3, 11]).

Infrastructures are being upgraded and server prices are coming down. Some production can be automated. It is therefore of increasing interest to see when and how video-on-demand is actually used when available.

Recent studies of the effectiveness of digital video-on-demand have focused on comparing distance learning via video on-demand to traditional classroom education [5, 6]. Our focus is on exploring usage patterns through detailed analysis of access logs and deriving implications for design of online presentations. For two years, videos have been available on the Microsoft corporate intranet for internal training and for viewing special presentations. This paper presents an analysis of over 33000 sessions of video access by more than 9000 people, enabling us to draw conclusions about designing for this important medium.¹ People are using the system, accessing some talks many months after the initial presentation. They use features that allow them to skip or browse. Most viewing sessions are brief.

This paper analyzes this large corpus in depth, both to confirm the usefulness of multimedia on-demand and to draw conclusions for those designing presentations and designing tools for authoring and delivering them. Of course, this environment is special in many respects, as is often true of early adopters. Care must be taken to consider these results in the larger emerging picture of web activity.

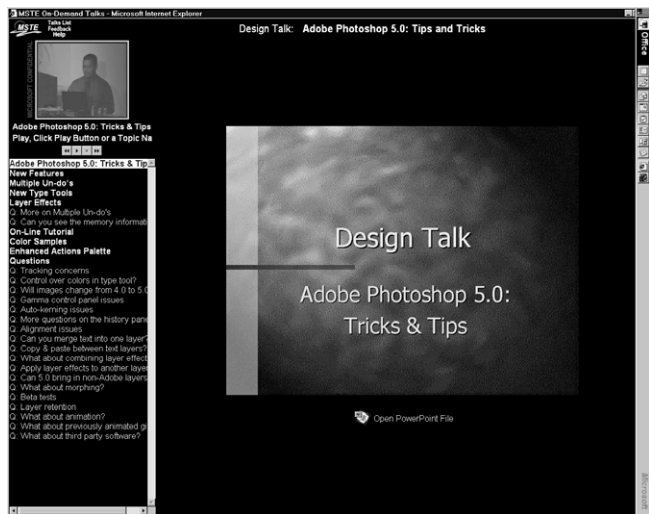


Figure 1. MSTE interface: video, slides, table of contents.

SYSTEM AND CONTENT

Microsoft Technical Education (MSTE) provides internal technical education to corporate employees. MSTE holds regular courses on software development, testing, and program management. They also produce company-wide

seminar series. Over two years ago, MSTE started digitizing videos of these talks and making them available online. MSTE provides both live *and* on-demand online access, but detailed logs of live access patterns are not kept.

MSTE client user interface

MSTE provides a web-based interface for accessing the talks. Employees can locate talks using several topic indices or keyword searches. Once located, a talk can be viewed by double-clicking on the title.

The client software to view the talk is also web-based (Figure 1). The three interface sub-frames are the video in the top-left frame, a table of contents (TOC) in the bottom-left frame, and the speaker's slides on the right.

The video is streamed using Microsoft NetShow software. The video window is quite small (176x144 pixels) to keep the bandwidth low (100Kbps). Users have standard speed and direction controls. The video camera is usually focused on the presenter. The slides are synchronized with the video: when the speaker moves to a new slide, the slide automatically flips. The TOC consists of one bullet for each slide and each audience question.² When a viewer clicks on a TOC bullet, the presentation jumps to the corresponding point in the talk. Although there is no noticeable delay before the slide appears, the video and audio take up to 15 seconds to appear due to buffering requirements.

Producing the talks with TOC and slide synchronization takes greater effort than just providing the video, but it can significantly enhance the ability of viewers to browse the talks and focus on relevant portions. Below we describe how these facilities are actually used.




Figure 2: New talks by month.

MSTE client logging

The on-line videos are usually placed on the server within a week of the presentation. The table of contents construction and the synchronization of the video and slides is done manually. From June 1997 to August 1999 about 3.6 talks per week have been made available, 367 talks in all (Figure 2). Over 80% are between 40 and 110 minutes long.

¹ A 2-page report based on 20% of the data appeared as [8].

² A TOC based on slide titles and questions can be constructed by someone unfamiliar with the content.

Detailed logs of viewers' watching patterns form the basis of this study. Every viewer interaction with TOC bullets or play, stop, and the other video controls generates an event. Each event is time stamped and logged, along with the viewer's ID, as a database record. We logged about 515,000 records from 33,160 viewing sessions.

OVERALL USAGE PATTERNS

Global usage patterns can indicate whether people continue to access presentations on-demand, how access is distributed across talks, and the 'shelf life' of a talk.

Sessions by month

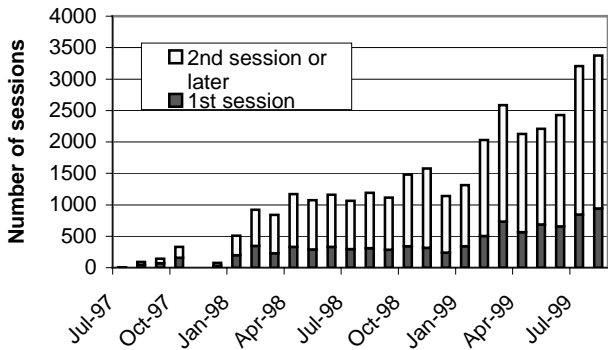


Figure 3: Sessions by month and viewer type. The apparent drop at the end of 1997 reflects an interruption in logging.

Video sessions increased steadily. (A *session* is a person accessing a given video at a single sitting. A person can spend multiple sessions on a talk.) Figure 3 shows these for first-time and repeat users. Most accesses are by repeat viewers, but there is a steady flow of first-time viewers. First-time use is likely to continue because of new hires, for whom these talks may provide especially valuable, convenient access to corporate knowledge.

Access to a talk over time

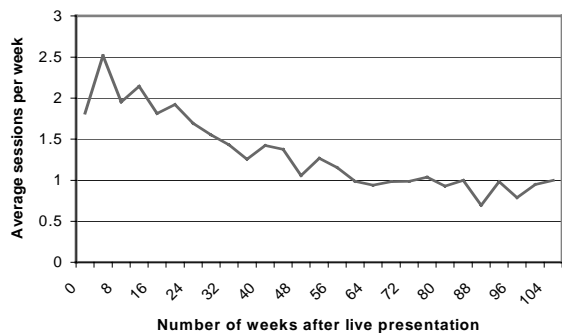


Figure 4: Distinct users accessing talks as a function of weeks since the presentation. (The number of talks contributing to the average declines along the X-axis, as most talks were online less than two years.)

How rapidly do accesses to a talk fall off over time? Given the storage space required by video files, this can indicate when a talk can be taken offline, freeing space.

Figure 4 shows these data averaged over all talks. Access peaks soon after the talk is put online and then declines, but the graph has a long tail. Even after two years, about 1 person per week accesses a talk on average (and this is in the fast changing software industry). Decisions about removing talks should be made on a talk-by-talk basis.

Audience size

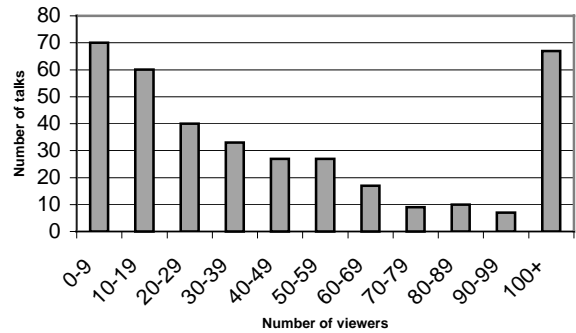


Figure 5: Number of talks as a function of online audience size.

An important measure of the value of putting a talk online is the number of people accessing it. Figure 5 shows the number of viewers for a talk on the X-axis and the number of talks on the Y-axis.

The median talk has had 34 online viewers. Sixty-seven of the 367 talks, mostly oriented toward software developers, have had 100 or more viewers. (The most popular talk to date is "XML in 180 Minutes," accessed by 1390 people with sessions averaging 33 minutes.)

The live lecture room attendance for MSTE presentations was tracked from November 1997 through August 1998. Average audience size was 83. Measured one year later, the average number of on-demand viewers for the same presentations was 84 (the median is 65, greater than the overall median of 34 because these talks have been available longer than average).

This level of on-demand viewing could justify the cost. The current process has two components: a few hours of production time for the camera crew to set up, record while multicasting online in real time, and tear down; and a few hours of post-production time to digitize and compress the video and to add slides and a table of contents.

The live network audience of employees watching from their desktops helps justify the production cost. Data are not routinely recorded, but as many as 9000 have viewed a presentation live via the intranet. Production and post-production have been streamlined, but camera management and index construction are still necessary.

ONLINE SESSION CHARACTERISTICS

This section focuses on viewer behavior while watching talks. For several of the analyses, viewers are categorized according to their previous access history, as extent of use may affect behavior. We examine how long people watch talks, to what extent they use the table of contents to skip

within a talk, what portions they watch, and so forth. These behaviors impact the ways that online talks could be structured or presented for more efficient online viewing.




Figure 6: Classification of users by number of talks watched. The actual numbers of users are on the top of each bar.

Classifying users by frequency of access

9095 distinct users have watched one or more MSTE talks online. This is about half of the company’s product development teams, which are the targeted audience for the MSTE series, and about 20% of the worldwide full-time and contingent staff, not all of whom have high bandwidth access to the corporate intranet.

As seen in Figure 6, a slight majority of viewers have so far accessed only one talk. The average is 2.7 talks. Frequent users, defined as those accessing 5 or more talks, comprise only 14% of the users but generated over 48% of the sessions. Given the growth in first-time use (Figure 3), a large number of one-time viewers does not imply an equally high attrition rate—almost 2000 first tried the system in the last two months surveyed, many of whom will use it again.




Figure 7: Number of sessions as a function of duration.

Session length

Figure 7 shows the number of sessions as a function of their duration. A majority, 57% are shorter than 5 minutes, and only 17% are longer than half an hour. This is in stark contrast to the behavior of live audiences. Most speakers would be horrified if half of the audience walked out within 5 minutes and only one fifth remained after half an hour.




Figure 8: Proportion of session watched for different groups.

To what degree does this reflect a first-time viewer, testing the system phenomenon? Figure 8 indicates that this is not the case. It shows the percentage of talks watched (rather than minutes watched) for viewers with different experience levels. Although a greater proportion of one-time user sessions are very short and a greater proportion of experienced user sessions are long, the differences are not great—there are many short sessions by experienced users and long sessions by first-timers.

This behavior shouldn’t come as a complete surprise. Online browsing is easy and requires less commitment than going to a room to see a presentation. There is little cost to discontinuing if a presentation is not engaging. There are more distractions in one’s office—phone, email, colleagues, reminders of other work. Furthermore, one can watch the summary or quickly skip through a talk and then resume other business. The effort to attend live talks is significantly greater, the time saved from skipping part of a lecture is proportionally smaller, and our desire not to offend the speaker or the host can further inhibit walking out.

Nevertheless, the strength of this phenomenon strongly suggests that presenters must get their message across quickly if they wish to reach online viewers. However, we need to know more. To what degree do viewers watch only the beginning? To what degree do they make use of the table of contents, including questions, created after the talk? Do they skim slides without watching the video?

Time coverage

Figure 9 shows the number of viewers for each minute of a typical talk. The overall effect is a saw-tooth pattern—the number of viewers tends to decrease as the talk progresses, and decreases within each slide, but the beginning of a slide often has more viewers than the final segment of the previous slide. This clearly reflects the use of the table of contents to skip to slides. The sharp rise at the beginning of the slide marked A is because it marks a new section, clearly evident in the TOC. The slide marked B shows no initial boost and a quick decline of viewers—it is labeled as a continuation of the previous slide. Viewers have little reason to skip to it; several who watched through to it apparently found little of interest and quit or skipped ahead.

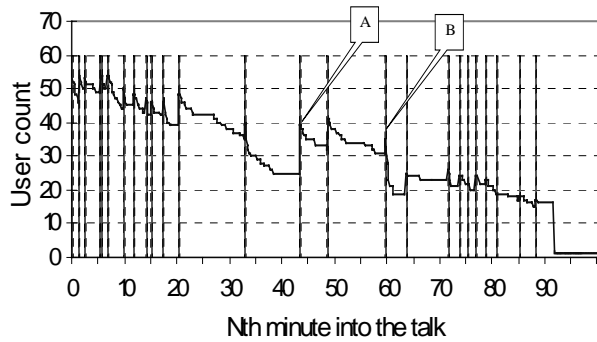


Figure 9: Viewers for each minute of a typical talk. Vertical lines denote presentation slide transitions.

In a presentation to be viewed online, it is critical to communicate key messages in the first few minutes. The same advice applies to each slide: Do not design a presentation based on experience with “live” audiences where most people stay regardless of the organization, delivery, or content of the talk. Web viewers are restless.

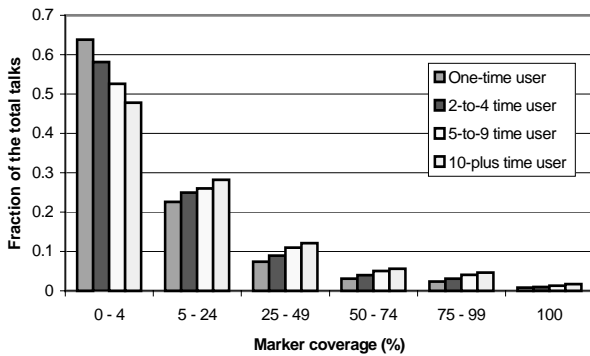


Figure 10: Fraction of sessions where a given percent of markers is covered, for different viewer experience levels.

Content coverage

Presentation ‘markers’ refer to table of contents entries, the links to slides and live audience questions. Viewers can access markers and the corresponding slide and video material by letting the video run or by using the table of contents entries to jump to them.

Figure 10 shows the fraction of sessions where a given percent of markers is accessed by viewers of that video. If a marker was not viewed for 15 seconds, it was not counted. The data are shown for the four levels of viewing experience. The pattern is similar to Figure 9: people often do not access many slides or markers, independent of their experience level.

First-time viewers tend to view few markers (86% of one-time users access fewer than 25% of markers). 10-plus-time users are more likely to cover more markers, with about 5% viewing 75% or more. And of course, slides at the beginning are more likely to be watched.




Figure 11: Percentage of sessions by number of jumps in each.

Linear versus non-linear access to video

How often do viewers use the index? They have the option of perusing a talk by flipping through the slides rather than watching the video. A user could spend 5 or 10 minutes viewing a talk yet go through most or all of the slides.

Here we examine the extent to which people jump and watch talks non-linearly. If this capability is not exploited, it may not be worthwhile to construct tables of contents.

Figure 11 shows over 60% of the sessions include at least one jump, although 83% have four or fewer.




Figure 12: Fraction of sessions with a given range of jumps.

Do many sessions have few jumps because they are short? Figure 12 compares data for sessions shorter than 5 minutes, 5 to 25 minutes, and longer than 25 minutes. Jump number does indeed correlate with session length, and we do find some long sessions that make no use of the index. However, it is very significant that even for talks of under 5 minutes, more often than not viewers make at least one jump. Thus, these very short sessions are not a completely casual access-and-quit, they involve active search.

Are frequent viewers more likely to use the index to access a video non-linearly? Figure 13 reveals no significant difference among groups.




Figure 13: Jump statistics among different viewer groups.

Table 1: Summary session statistics viewers of MSTE talks.

	One-time users	2-to-4 time users	5-to-9 time users	10-plus time users
Number of viewers in category	4786	3095	879	335
Average Session Length (min)	14.1	14.4	15.3	16.9
Average # of Sessions per Talk	1.32	1.36	1.46	1.59
Average # of Jumps per Talk	3.08	2.83	2.67	2.66
Average Time per Talk (min)	18.7	19.5	22.3	26.8
Total Sessions	6333	10919	8023	7885
Total Talks	4786	8041	5508	4966

Summary Data Based on User Characteristics

Table 1 presents a summary of session statistics for viewers of different experience levels. Some trends: i) The average session length increases with frequency of viewing. ii) The average number of sessions per talk (i.e., the same talk is viewed by the same person in multiple sittings rather than a single sitting) increased from 1.32 to 1.59 from one-time to 10+ users. Combining these two results, the 10+ users average about 27 minutes on each talk they view versus about 19 minutes for first-time viewers. iii) The average number of jumps decreases slightly with experience.

The behavior of experienced viewers differs from that of first-time viewers. Not every first-time viewer becomes a 10-time viewer, so a question arises: Did the profile of these 10-time viewers resemble the first-time viewers when they began, or do people who will continue to come back behave differently from the start?

In Table 2 we examine the evolution over time of the behavior of the 335 10+ viewers: What was their first-time behavior, second to fourth time use, and so on. (Note that only the first columns of the two tables are directly comparable; for example, the second column of Table 1 includes the first and subsequent talks of 2-time, 3-time, and 4-time viewers, whereas the second column of table 2 contains the second through fourth talks of 4-time users.)

We see that the people who will turn out to be frequent viewers are more likely to spend multiple sessions with

their first video and average 46% more time on it than first-time viewers overall. The average time they spend on a talk stays relatively constant at about 28 minutes, but they shift to fewer, longer sessions and somewhat less jumping. This could reflect more efficient use of the system.

Table 2: Evolution of behavior of 335 10-plus time viewers.

	1st	2 nd -4 th	5 th -9 th	10 th +
Average Session Length (min)	14.3	15.7	16.5	18.5
Average # of Sessions per Talk	1.90	1.79	1.49	1.51
Average # of Jumps per Talk	2.95	2.60	2.70	2.58
Average Time per Talk (min)	27.2	28.1	24.6	27.9
Total Sessions	637	1802	2492	2954
Total Talks	335	1005	1675	1951

DISCUSSION

In the high-technology environment we examined, on-demand video presentations are accessed by a substantial fraction of employees. Given the pace of technology change, many organizations will soon have the infrastructure and server capacity to host on-demand video. Small organizations may not generate as much content internally, but all organizations need to learn; third parties are already creating content and making it available online. This study confirms that people find informational video useful. On-demand viewing of talks designed for live presentation now exceeds local attendance, motivating fresh thinking about the design of such talks.

Viewers of on-demand presentations behave differently than the audience in the room. They usually spend much less time on a presentation, but they do not simply watch for a few minutes and then stop. With a table of contents or index into the presentation, even short sessions are frequently accompanied by one or more jumps. First time viewers are almost as likely to use the index as experienced viewers. In addition, one third of the time first time viewers return to view a video more than once.

These results parallel the findings of Morkes and Nielsen [13, 15] for reading text on the web: This is browsing and skimming more than viewing (or reading). But unlike most web page designers, a person preparing a presentation must consider the needs of traditional viewers (in the live, local audience) and the needs of the on-demand audience. They may also have a live but remote audience. In some cases a presentation design that works for one also works for the others; in other cases there are tradeoffs.

Because on-demand attention drops off steadily and relatively sharply with time, speakers should emphasize key points immediately in the talk. This is a good policy for local audiences, and even more significant for on-demand audiences, who are trying to determine if a talk is germane to them. Saving a surprise result to be a “punch line” may work in the presentation room, but is ill-advised for on-demand. The person in the live audience has usually

committed the full block of time. The on-demand viewer may not intend to watch the entire talk even if it is germane.

Similarly, with attention highest at the beginning of each slide, key points should be emphasized immediately following slide transitions. Nielsen and Morke recommend limiting each web paragraph to a single idea, because readers are skimming. Similarly, presenters should consider limiting each slide to one key point. The second point on a slide is more likely to be missed. Of course, if taken to an extreme, slide-switching could distract the local audience.

Speakers are often advised to increase legibility by minimizing text in slides. In fact, they are encouraged to shift the focus to themselves and their words by minimizing slide use altogether. On-demand viewers do not have a legibility problem, they use slide content to browse, and get less from the small video image of the speaker. So providing them with more text is a valid design goal that creates a conflict with designing for the live audience.

Can this be reconciled? With care, more detail could be added to slides after a presentation, if care is taken not to lose synchrony with what was said. But changing slide content appreciably for on-demand viewing will create a mismatch between the context in which words were spoken and that encountered by the on-demand viewer

A major distinction between local and on-demand presentations is the significance of slide titles. A local audience has the prior context of a talk to help interpret slide titles; viewers may look at the first bullet without even reading the title. A slide title can be a humorous transition or use a term introduced earlier in the talk and be fully understood. But slide titles that form a table of contents should fully describe the contents and be understandable outside the presentation context. This is similar to the importance of web page titles and headers. For example, the header for this section, "Discussion," is OK for someone reading the paper, but a person browsing might prefer "Implications of Use & Guidelines for Design."

Going a step further, slide titles should reveal as much of the structure of a talk as possible. Major sections or topic shifts that are clearly reflected in slide titles will benefit on-demand viewers. Talk titles themselves should be carefully considered with the browsing viewer in mind. The MSTE set included three series, whose segments are listed as "Module 1," "Module 2," etc. More specific titles could be particularly useful to potential online viewers.

References to related work are always useful, but URLs (active links if possible) could particularly appeal to on-line viewers (some of whom suggested this feature).

The results also guide those building tools to support the viewing and authoring of digitized presentations. Greater support for skimming and browsing is possible, including automatic generation of multimedia summaries [8]. Mixed-initiative authoring tools are possible. Presenters invest time preparing a talk for a live audience; when more than half of the total audience could be on-demand, they may be happy

to contribute to post-processing the talk. Even a few minutes could contribute significantly to preparing better indices or summaries: presenters could indicate the most and least important slides, re-title some to be clearer, identify important bullet items, locate relevant online information, and so forth.

An ambitious approach to improving on-demand viewing is software that automates the construction and presentation of a hierarchical index that reveals the structure of a talk [1, 12, 14, 17].

Compression, skimming, annotation: increasing the value of on-demand viewing

Live viewing has two major benefits: immediate reception of information and the potential for interaction with the speaker and other audience members.

Viewing an event 'live' can have a strong psychological effect – it is different than seeing a recorded broadcast. Perhaps it is the possibility that something remarkable will happen in a live event that we know is not in the rebroadcast because we would have heard about it. This will not change, but emerging technologies are creating compensatory benefits that could greatly increase our interest in watching informational video on demand.

Software that time-compresses audio and video while maintaining audio pitch enables viewers to get more information in less time [2, 7, 16]. People can comfortably view videos at speeds 1.5 to 2 times normal: an hour lecture in 30 to 45 minutes. To enable yet more efficient viewing by supporting the skimming or browsing we observed, we are exploring the automatic extraction of summaries, segments of audio and video that constitute presentation highlights [8, 9]. Segment selection can be based on slide transitions, speech intonation, or "collaborative filtering," steering viewers toward portions of a talk where previous viewers spent more time. Tools that allow speakers to help in extracting such highlights could improve them further.

Interacting and collaborating around digitized multimedia can be facilitated by technologies that allow people to easily create and share annotations that are anchored to specific video segments [3, 4]. Such technologies can only make on-demand viewing more appealing.

CONCLUDING REMARKS

The last few years have seen a rapid maturation of basic streaming media technologies. A key application is the online, on-demand availability of informational talks. Based on detailed usage logs, we have suggested how such access can be made more useful at small cost.

The corporate talks that we studied were often viewed months after being given. They were well-attended in person; even more people viewed them on-demand. As storage costs decline, the effort to make them available is more easily justified, motivating ways to enhance on-demand viewing experience.

Viewers approach archived multimedia presentations by browsing and skimming. We have outlined ways to support

this, if possible without detracting from the experience of the live, local audience, yet realizing that design decisions can differentially affect each set of viewers.

A third audience is the live but remote viewers of presentations, notably those viewing over a network from their offices. These viewers cannot benefit from post-processing but can be provided with means of interacting with the speaker or each other [11] and can be given pointers to ancillary material that could be viewed online during a presentation. The compression technologies mentioned in the last section potentially blur the live and on-demand distinction: a viewer who joins late could review the beginning of a talk at a compressed speed and 'catch up.'

A critical aspect of research in the uses of multimedia will be to closely observe experiences with different technologies in different domains. Making effective use of new technologies is challenging. Efforts often take unexpected turns. Understanding and sharing initial experiences is critical to making progress.

We are at the beginning of the exploration of interactive online multimedia. The path is unpredictable. If presenters follow the guidelines suggested here and design talks considering both live and on-demand viewing, it could change the experience for all audiences.

REFERENCES

- [1] Abowd, G.D., Brotherton, J. & Bhalodia, J., 1998. Classroom 2000: A system for capturing and accessing multimedia classroom experiences. *CHI'98 Extended Abstracts*, 20-21. ACM.
- [2] Arons, B., 1997. SpeechSkimmer: A system for interactively skimming recorded speech. *ACM Transactions on Computer Human Interaction*, 4, 1, 3-38.
- [3] Barger, D., Gupta, A., Grudin, J. & Sanocki, E., 1999. Annotations for streaming video on the web: System design and usage studies. *Proc. WWW8 World Wide Web Conference*, 61-75.
- [4] Barger, D., Gupta, A., Grudin, J., Sanocki, E. & Li, F., in press. Asynchronous collaboration around multimedia and its application to on-demand training. To appear in *Proc. HICSS-34*.
- [5] Belanger, P.C. & Clement, S., 1995. Using video-on-demand for educational purposes: Observations from a three-month experiment. *Canadian Journal of Educational Communication*, 24, 1, 61-83.
- [6] Branch, P. & Durran, J., 1996. PC based video on demand trials. In *Learning Technologies: Prospects and pathways*. Selected papers from *EdTech '96 Biennial Conference of the Australian Society for Educational Technology*.
- [7] Harrigan, K.A., 1996. Just noticeable difference and effects of searching of user-controlled time-compressed digital-video. Ph.D. Thesis, University of Toronto.
- [8] He, L., Sanocki, E., Gupta, A. & Grudin, J., 1999. Auto-summarization of audio-video presentations. In *Proc. Multimedia'99*, 489-498. ACM.
- [9] He, L., Sanocki, E., Gupta, A. & Grudin, J., 2000. Comparing presentation summaries: Slides vs. reading vs. listening. *Proc. CHI 2000*, 177-184. ACM.
- [10] He, L. Gupta, A. & Grudin, J., 1999. Design lessons from deployment of on-demand video. *CHI'99 Extended Abstracts*, 276-277. ACM.
- [11] Jancke, G., Grudin, J. & Gupta, A., 1999. Presenting to local and remote audiences: Design and use of the TELEP system. *Proc. CHI 2000*, 384-391.
- [12] Ju, S.X., Black, M.J., Minneman, S. & Kimber, D., 1997. Analysis of gesture and action in technical talks for video indexing. *Proc. CVPR'97*, 595-601. IEEE.
- [13] Morkes, J. & Nielson, J., 1997. Concise, SCANNABLE, and objective: How to write for the Web. <http://www.useit.com/papers/webwriting/writing.html>
- [14] Mukhopadhyay, S. & Smith, B., 2000. Passive capture and structuring of lectures. *Proc. Multimedia'99*, 477-487. ACM.
- [15] Nielsen, J., 1997. Guidelines for multimedia on the Web. <http://www.useit.com/alertbox/9512.html>
- [16] Omoigui, N., He, L. Gupta, A., Grudin, J., & Sanocki, E., 1999. Time-compression: System concerns, usage, and benefits. *Proc. CHI'99*, 136-143. ACM.
- [17] Uchihashi, S., Foote, J., Girgensohn A., & Boreczky J., 1999. Video Manga: generating semantically meaningful video summaries. *Proc. Multimedia'99*, 383-392. ACM.