

Framework for effective use of multiple displays

Qiong Liu¹, Don Kimber¹, Frank Zhao¹, Jeffrey Huang²

¹FX Palo Alto Laboratory, 3400 Hillview Ave. Bldg. 4, Palo Alto, CA 94304

²Harvard Design School, 48 Quincy Street, Cambridge, MA 02138

Abstract

Meeting environments, such as conference rooms, executive briefing centers, and exhibition spaces, are now commonly equipped with multiple displays, and will become increasingly display-rich in the future. Existing authoring / presentation tools such as PowerPoint, however, provide little support for effective utilization of multiple displays. Even using advanced multi-display enabled multimedia presentation tools, the task of assigning material to displays is tedious and distracts presenters from focusing on content.

This paper describes a framework for automatically assigning presentation material to displays, based on a model of the quality of views of audience members. The framework is based on a model of visual fidelity which takes into account presentation content, audience members' locations, the limited resolution of human eyes, and display location, orientation, size, resolution, and frame rate. The model can be used to determine presentation material placement based on average or worst case audience member view quality, and to warn about material that would be illegible.

By integrating this framework with a previous system for multi-display presentation [PreAuthor, others], we created a tool that accepts PowerPoint and/or other media input files, and automatically generates a layout of material onto displays for each state of the presentation. The tool also provides an interface allowing the presenter to modify the automatically generated layout before or during the actual presentation. This paper discusses the framework, possible application scenarios, examples of the system behavior, and our experience with system use.

1. Introduction

In modern society, people are surrounded by arrays of multiple, simultaneous images. Using multiple displays to convey information can be traced back to early television studios packed with walls of displays or the so-called Situation Rooms used during World War II. Other than specific usage of multiple displays in early days, the Eameses pioneered this technique in public shows. In the 1959 American exhibition in Moscow, the Eameses showed their film "Glimpse of the USA" with seven twenty-by-thirty-foot screens suspended within a vast (250 feet in diameter) dome. The show turned out to be a great success [13]. Besides the American exhibition, the Eameses and their followers designed many other impressive presentations and shows using multiple displays [14].

Even though multiple displays can facilitate a presenter to convey sophisticated knowledge to audience members and we are surrounded by multiple displays in many public places, existing presentation tools, such as PowerPoint, provide no support for multiple displays. To enable better use of media-rich environments, we designed EPIC, a tool for authoring and replaying presentations on arbitrary device configurations. EPIC complements, but does not replace tools used to author specific media. It can organize media prepared for simple devices and synchronously present them in one or more multimedia venues. For example, the EPIC system can import a conventional PowerPoint file and re-author it for effective presentation on multiple displays. Our prototype supports arbitrary configurations of displays, printers, speakers and room lights.

EPIC, which stands for Environment Picking Image Canvas, uses live or static images of the presentation environment as a graphical user interface (GUI). This allows users to visually select and control presentation devices, even in remote locations. Users may drag slides or media files onto a visual representation of the intended device. For example, a user can drag a slide thumbnail on top of any of the displays visible in Figure 1 to show a slide on that display. Additionally, EPIC supports virtual presentation previews. By displaying the selected slides overlaid on the room image, video, or 3D model, users can see what their presentation will look like in advance. Figure 1 shows the EPIC interface, which includes the active image/video canvas, a hyper-slide pool, a zoom pane to see slide details, and a timeline for arranging presentation slides on multiple devices. The GUI also has four tool bars, (marked 1-4), for control device definition, file manipulation, presentation control, and timeline manipulation respectively.

The device-state table is a timeline interface that maps ‘hyperslides’ to room devices according to the presentation state. (We use the term ‘hyperslide’ for anything that can be rendered on a device, such as a presentation graphic, an audio or video file, a print document, or a room control action such as “dim the lights.”) In the table, each row shows the hyperslides associated with a particular device, and each column shows the hyperslides scheduled at a particular event time. To schedule a hyperslide on device X at time Y, the user drags the hyperslide thumbnail to the location at row X and column Y. During the presentation, the user advances the timeline using the PowerPoint conventions of a down arrow or mouse click, or a ‘slide advance’ button. All hyperslides scheduled at a particular event time are then sent to their respective devices, allowing synchronous changes across heterogeneous displays.

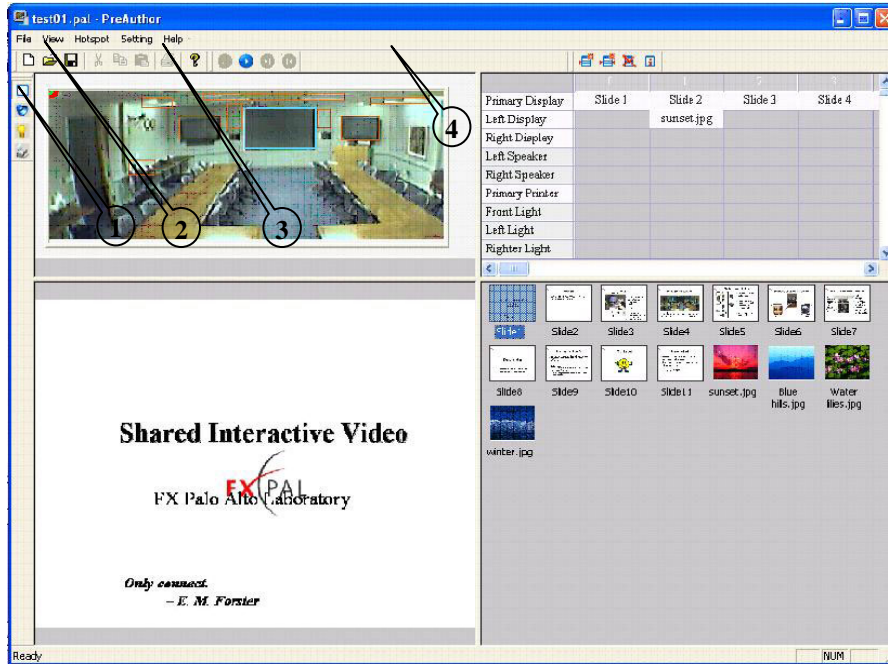


Figure 1. Main User Interface of EPIC (a)Top-left: An environmental image/video canvas depicts the device configuration of a venue. (b)Top-right: A device-state table reveals hyper-slides’ relations for a unified presentation. (c)Bottom-left: A zoom pane for checking details of a user’s selection. (d)Bottom-right: A hyper-slide pool shows thumbnails of all hyper-slides.

The remainder of this paper is organized as follows. The next section describes previous work and systems related to our authoring tool. Section 3 gives a system overview. Section 4 covers the EPIC user interface in depth. Our detailed strategy for computer assisted presentation content management is given in section 5. Section 6 and 7 are dedicated to EPIC application scenarios and examples respectively. Concluding remarks and future work are given in section 8.

2. Related Work

Many commercial products are available for presentation authoring. Microsoft PowerPoint, Freelance [3], and OpenOffice [5] are typical examples. They organize presentation contents as a single slide sequence, and support text editing, figure drawing, audio/video embedding and animation for each slide. These tools support presentation authoring only for a single display and stereo audio channel. Associating slides with devices is trivial in this case, so computer assistance is unnecessary.

ConferenceXP [2] can handle multiple devices during presentation. This software [7, 8] distributes the same content to all displays for audience members, and does not support distribution of different contents to various devices. Therefore, this software won’t support arbitrary device configurations. For example, showing the same content on multiple large displays is less informative than individually controlling each. Additionally, Conference XP organizes presentations as a single sequence of marker-enhanced slides. Therefore, computer assisted device-content association is not necessary for this software.

Systems described in [4, 6] can support online manipulation of multiple displays. However, they are not designed for presentation authoring, and therefore difficult to use for presentation preparation. Moreover, they are limited to displays, and cannot support other devices, such as multiple loudspeakers or lights that may be available.

EPIC and our previous work [1] tackled different presentation problems. In [1] the authors focused on interactions with multiple devices during playback. EPIC, however, emphasizes the authoring interface that can handle arbitrary device configurations. EPIC also includes some techniques described in [1] for online content manipulation during presentation.

Compared with an earlier system named PreAuthor [9], EPIC provides computer assisted slide management functionality. With this novel functionality, EPIC can assign slides to various devices according to users' guidelines, and presentation venue configurations. This is much more convenient than the manual slide-device association implemented in PreAuthor. Moreover, EPIC's real-time content manipulation allows presenters to extemporaneously modify the presentation. For example, the presenter may display additional slides or annotate existing slides in response to audience questions.

3. System Overview

EPIC supports authoring and playback of synchronized presentation sequences for arbitrary combinations and placement of devices. EPIC is not a replacement for the many media editing tools [3, 5] in widespread use; instead, it is designed to organize pre-edited slides, audio files, images, and movies for multiple devices connected to the same network. It is a higher-level tool for composing results of various single-channel media editors into a unified presentation across multiple devices.

Before getting into design details, it is helpful to introduce some terminologies used by EPIC.

Venue: a venue is a setting in which a presentation occurs. It may be a single room, or distributed, like the multiple locations in a teleconference.

Hyperslide: The hyperslide, or h-slide, is a basic presentation unit of EPIC. A hyperslide is an input source or an object that can be 'rendered' by a device. Examples include a regular PowerPoint slide, an image, a video clip, an audio segment, a web page, streaming audio/video from a live microphone or camera, a printable file, or even a light control command (on/off/dim).

Media Device: a *media output device* (e.g. projectors, printers, or loudspeaker systems) renders hyperslides. A *media input device* such as a microphone or camera provides input to a hyperslide stream.




Figure 2. An Environmental Image/Video Canvas

Device Portals: The EPIC interface controls devices using a visual metaphor. Devices are selected and accessed through their image in an active window containing an image or video, or rendering of the venue. A *device portal* is the active region in the window associated with a device. Device portals are indicated by highlighted boxes in the active image. For example, Figure 2 shows four device portals corresponding to a printer (left) and three displays. These portals are "active" for drag-and-drop operations; dragging a hyperslide to a portal will schedule or render it on the associated display.

Channel: a channel is an 'abstract device' to which h-slides can be associated. A channel can be mapped to one or more physical devices. For example, the *primary display* channel is typically rendered on the most prominent display(s) in a venue. A *notes channel* may 'broadcast' some h-slides to devices such as audience members' laptop displays. A *video channel* may be mapped to both visual and audio output devices, typically a visual display and a loudspeaker.

Device-State Table (DST): A DST is used to organize hyperslides across time and devices. It is also useful for revealing h-slides' relations on a display. Each row of DST corresponds to an available channel, while each column of

DST corresponds to a state in the presentation. When manually authoring an EPIC presentation, the user schedules hyperslides by dragging thumbnails to their desired locations in the DST. EPIC also supports automatic authoring, where the system generates an appropriate DST for a given input. Automatic authoring is a major topic of this paper and we present methods for determining proper DST slots for all h-slides.

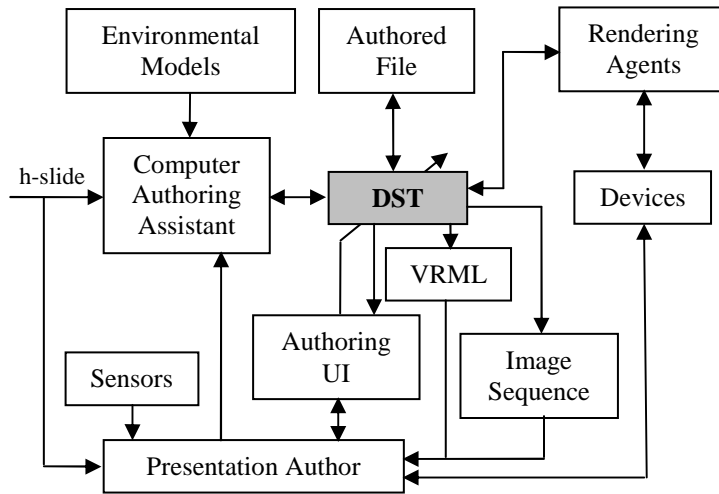


Figure 3. System Architecture

Figure 3 shows the overall system architecture of EPIC. A presentation author interacts with the DST using the authoring user interface, which allows EPIC presentations to be loaded, authored, saved, and vividly previewed with an image sequence simulation, or a 3D VRML simulation. For presentation, the author may select local or remote playback, and can modify the presentation in real time by dragging h-slides onto device portals. The presenter uses standard keyboard controls such as arrows to change the presentation states. Depending on the venue configuration, various sensors, such as cameras, clocks, and touch screens, can also be used to control the presentation progress.

To simplify authoring, EPIC can automatically lay out a presentation using rules and author preferences, device and audience location information, and the presentation contents. To improve automatic authoring performance, it is able to capture DST statistics for future reference.

A multiple-device EPIC presentation works through network unicast. Remote agents run on every host to which one or more devices are connected. These agents listen on a pre-defined IP port. The EPIC application sends h-slide contents to a device agent, which then spawns an application to render the h-slide on the corresponding device. Broadcast channels, such as the notes channel, are implemented by making h-slides associated with the channel available via HTTP at channel-specific URLs. Note that the EPIC application can be run anywhere, given network connectivity, so that presentations may be delivered in remote locations using the same interface. We have successfully used the EPIC system in trans-Pacific teleconferences for multi-display presentations to our remote colleagues.

4. Authoring and Playback UI

As shown in Figure 1, the EPIC UI consists of four panes. In the top left pane, an active canvas displays an image of the venue and any portals associated with venue devices. The venue image may be a static image, a live video panorama, or a rendering of a 3-D VRML model. A static image is small and easy to store, while the video image has the advantage of showing the current state of the display as well as any people in the venue. The VRML model has the advantage that presentations may be previewed from any viewpoint in the venue. The bottom right of the EPIC window shows thumbnails of the h-slides available for presentation. Selecting a thumbnail shows the slide in more detail in the bottom left pane; thumbnails can be dragged to arrange them in the DST at top right. The DST pane allows users to see and to specify which h-slides are rendered on which device at each state of the presentation. In Figure 1, slide 1 has been selected, so its contents are shown in the zoom pane.

The UI also contains four tool bars for authoring and playback control (marked 1-4 in Figure 1). Tool bar 1 has buttons for each device type – screen, speaker, room control, and printer – and is used when defining a portal. Tool bar 2 is for file manipulation such as opening, saving, printing, etc. Tool bar 3 controls presentation playback such as start and slide advance. Tool bar 4 includes functions for DST manipulation.

4.1. Portal Configuration

Before the EPIC system can be used to control remote devices, they must be associated with portals and output channels. Usually, this need be done only once for a venue; the casual user will use existing portal definitions. When one adds a device or configures a new venue, portals must be defined for each device. Each portal includes six properties: name, type, computer host, connection port, and portal location and size in the active image. To define the latter two, the user may draw a bounding box using the mouse. After defining the portal location, a dialog box prompts the user for additional portal properties.

Portal definitions for all devices in a venue can be saved in a template file with an environmental model or picture. For each meeting room configuration, the template file only needs to be created once. After the creation of a meeting room configuration file, every user can use this template to assist presentation authoring. The template file can be sent to remote users to prepare a presentation in advance, or to remotely control the local room in real time.

4.2. Presentation Authoring

To author a presentation, the user arranges h-slides in the DST by dragging and dropping; the usual editing commands such as copy, paste, insert, and delete are also supported. An individual thumbnail representation is generated for every h-slide, and EPIC allows the user to manipulate them in various ways. In authoring mode, a user can drag an h-slide thumbnail onto a portal to indicate the h-slide should be displayed on the device associated with that portal. After the drag and drop action, the h-slide will appear in the DST based on that device and the current state. (In presentation mode, the h-slide is actually displayed by the indicated device). H-slides located in various DST cells can be moved and rearranged. The user may also double-click on an h-slide thumbnail to launch a tool for editing that type of h-slide (e.g. PowerPoint for a PPT slide.)

4.3. Presentation Preview and Playback

EPIC provides both a preview mode and a playback mode. In the preview mode, the actual room devices are not controlled. Rather, a preview of the presentation is rendered in the image canvas. By displaying h-slide thumbnails in the portal bounding boxes (see Figure 4), the user can navigate the presentation using normal controls, and see a visual indication of what h-slides are rendered on each visible device. The image canvas can contain a rendering of a VRML model of the venue, if available. In this case, h-slide thumbnails are rendered in 3-D on the display surfaces. Because the rendering viewpoint can be selected, the user can preview the presentation from any point in the venue. Figure 5 shows an example of a VRML output, including a pop-up menu for selecting the viewpoint. Users may also zoom in and out of the rendering to see details or overviews of the presentation.




Figure 4. A Rendered Image for Authoring Result Preview.

During playback, EPIC sends synchronized commands to multiple networked devices. Unlike familiar presentation tools that only activate a slide change on one display with each key-press, EPIC sends a set of synchronized media rendering commands to multiple devices with each key-press.

Presentation venues may be distributed across multiple locations for teleconferences or distance learning. The EPIC image canvas can show live video from a remote conference room, and users can see zoomed-in camera views in the detail pane. This feature is useful for giving a presentation to a remote site. Details of controlling and sharing remote cameras are described in [12].



Figure 5. A VRML Output for Authoring Result Preview.

Besides playing back previously authored presentations, EPIC also supports real-time slide presentation. In presentation mode, a user may drag an h-slide from a desktop or the slide pool to a portal for immediate rendering on the device associated with that portal. This updates the DST, and on exit the user is asked whether to save the modified DST.

5. Computer Assisted Authoring

Arranging a large number of h-slides for multiple displays is a tedious task. To assist the user, EPIC offers automatic algorithms for creating multi-display presentations. When presentation material (typically from PowerPoint) is imported into EPIC, simple rules determine which h-slides are displayed at each presentation step. These are generated automatically, but can be modified by the author. The system then automatically finds the ‘best’ mapping from h-slides to devices for each state. Finally the author may do some ‘fine tuning’ by overriding any of the automatic assignments.

To allow more control over the initial mapping of imported content, rules may be specified to help assign playback devices or transitions for h-slides. Examples of these are ‘current slide on primary-display’, ‘notes on audience-PDA-and-laptop-displays’, ‘outline on left-display’, or ‘display previous-slides’, etc. These rules constrain the choices of mapping for the DST.

When preparing a presentation, a basic goal is to let audience members perceive presentation materials as clearly as possible. EPIC models the quality of view available to audience members to find the best mapping from h-slides to devices, subject to constraints and rules described above. (We discuss here the case of visual media, but similar analysis could be given for audio.) A user’s perception of an h-slide depends on many parameters, such as distance between a user and a display, the display’s size and resolution, etc. The goal of our management algorithm is to let audience members perceive displayed image, video or text with the minimum distortion.




Figure 6. Our Model of the Information Flow to the User’s Visual Perception System

A user’s view of a visual signal may be modeled with the block diagram shown in Figure 6. By using u , v , and t to represent horizontal coordinates, vertical coordinates, and time respectively, an ideal signal, $f(u, v, t)$, generated by an

authoring system passes through a ‘‘Display Filter’’ and a ‘‘Space Filter’’ before it becomes $\hat{f}(u, v, t)$ as perceived by the user. In this diagram, the ‘Display Filter’ is used to model the limited resolution of a display. It may be considered as a band-limited filter whose horizontal cut-off frequency ωd_h and vertical cut-off frequency ωd_v are equal to one-half of the horizontal and vertical display resolutions respectively. The ‘Space Filter’ is used to model the spatial relationship between the user and a display patch, and the limited resolution of the user’s eyes. It may also be considered as a band-limited filter with a cutoff frequency equal to one-half of the user eye’s resolution. Conceptually, $\hat{f}(u, v, t)$ may be thought of as the best possible reconstruction of the signal $f(u, v, t)$ from a camera at the position of the user’s eye, and with resolution comparable to the eye. Assuming the spatial resolution of a user’s vision is homogeneous in various directions, we may denote the spatial cut-off frequency by ω_s , and the temporal cut-off frequency by ω_t .




Figure 7. Scan Direction of a Display

Conceptually, an audience member’s location is considered as a point (x, y, z) in world Cartesian coordinates, and a point on a display has parameters $(x_1, y_1, z_1, \varphi_1, \theta_1, r_{h1}, r_{v1}, r_{t1})$, where (x_1, y_1, z_1) reflects the position of the point, (φ_1, θ_1) gives us the scan direction of the display like that shown in Figure 7, r_{h1} , r_{v1} , and r_{t1} are horizontal resolution, vertical resolution, and frame rate of the display respectively. Denoting $X = (x, y, z)^T$, the perception scaling factor, α , of a horizontal line may be approximated with

$$\alpha = \sqrt{1 - \frac{[(X_1 - X) \cdot \varphi_1]^2}{|X_1 - X| \cdot |\varphi_1|}} \Big/ |X_1 - X|$$

Similarly, the perception scaling factor, β , of a vertical line can be approximated with

$$\beta = \sqrt{1 - \frac{[(X_1 - X) \cdot \theta_1]^2}{|X_1 - X| \cdot |\theta_1|}} \Big/ |X_1 - X|$$

By using F to represent signals in the spatial frequency domain and assuming displays and human eyes act as band-limited filters, the signal relations in the model may be described with following equations:

$$F_d(\omega_u, \omega_v, \omega_t) = \begin{cases} F(\omega_u, \omega_v, \omega_t) & * \\ 0 & \text{Otherwise} \end{cases}$$

$$* \quad \omega_u \leq \frac{r_{h1}}{2}, \omega_v \leq \frac{r_{v1}}{2}, \omega_t \leq \frac{r_{t1}}{2}$$

$$\hat{F}(\omega_u, \omega_v, \omega_t) = \begin{cases} F_d(\omega_u, \omega_v, \omega_t) & ** \\ 0 & \text{Otherwise} \end{cases}$$

$$** \quad \omega_u \leq \alpha \omega_s, \quad \omega_v \leq \beta \omega_s, \quad \omega_t \leq \omega_{st}$$

With these equations in mind, the content distortion, D_c , of a perceived visual signal may be estimated with:

$$D_c \propto \iiint_{\substack{\omega_u > \min(\alpha \cdot \omega_v, \frac{r_{h1}}{2}) \\ \omega_v > \min(\beta \cdot \omega_u, \frac{r_{v1}}{2}) \\ \omega_t > \min(\omega_u, \frac{r_{t1}}{2}) \\ \text{from } t \text{ to } t+T}} |F(\omega_u, \omega_v, \omega_t)|^2 d\omega_u d\omega_v d\omega_t .$$

D_c may be used to measure the visual distortion when a slide is correctly assigned to a display. When a computer automates slide assignment, its choices may differ from the desired choices of the author. The corresponding ‘loss’ when a slide is incorrectly assigned to a display could be modeled in various ways, but for simplicity we take it to be

$$D_{inc} \propto \iiint_{\substack{\omega_u \geq 0 \\ \omega_v \geq 0 \\ \omega_t \geq 0 \\ \text{Over } T}} |F(\omega_u, \omega_v, \omega_t)|^2 d\omega_u d\omega_v d\omega_t ,$$

where F is the spectra of the slide that was displayed incorrectly.

Assume $\{R_i\}$ is a set of non-overlapping small regions on a display, T is a short time period, $p_t(R_i | O)$ is the probability of correctly assigning R_i conditioned on O , where O is a conditional state corresponding to context and possibly environmental observations. O can include features from text on a slide, the state of an h-slide, or textures within an image. The overall information loss of assigning a visual object to a display may be estimated with

$$D = \sum_i \{p_t(R_i | O) \cdot D_{c,i} + [1 - p_t(R_i | O)] \cdot D_{inc,i}\}$$

In the above equation, it is assumed that the probability does not change during a relatively long period. This probability may be estimated with

$$p_t(R_i | O) = \begin{cases} 0 & \text{Against Guidance} \\ \frac{p_t(O | R_i) \cdot p_t(R_i)}{p_t(O)} & \text{Otherwise} \end{cases}$$

where the users’ guidance to the systems may be provided to the system as rules, such as, ‘on primary-display’, ‘on left-display’, ‘on audience-PDA&Laptop-displays’, ‘on all-displays’, ‘three h-slides in every state’, or ‘one-to-one association enforced’ etc. The probability of satisfying an h-slide arrangement in a region and the probability of using a region may be estimated based on the system’s past experience. Since users’ preferences are evolving processes that may be affected by good/bad presentations these users attended, the above probability estimations should be online learning processes to adapt to these evolutions. That is why we put t as a footnote for all probability estimations.

To achieve the best authoring result, our h-slide-display association strategy is to minimize the overall distortion D . Assume $\{S_i\}$ is a list of h-slides, $\{Device_i\}$ is a list of devices corresponding to the list $\{S_i\}$. The optimal device assignment list $\{Device_i\}_o$ may be described with

$$\{Device_i\}_o = \arg \min_{\{Device_i\}} (D) .$$

With this h-slide-device association strategy, the computer system can support a range of options from untended automatic to full manual device-h-slide association. This strategy is also consistent with our intuitions on slide assignment. For example, we prefer using large, high resolution displays to show our slides; we prefer allocating large, high-resolution display for images that have more details; we prefer using displays closer to all audience members; we prefer giving users handouts when the display size and resolution is not enough for us to show details.

Figure 8 shows the top view of a typical wide conference room that we may find in many conference hotels and institutions. If we do not limit the number of states for a presentation, and do not give the system any guidance for h-slide arrangement, the computer algorithm may automatically show every slide on all displays for a better viewing result. This is consistent with the common practice performed in this kind of rooms.




Figure 8. A Wide Conference Room with Multiple Displays

If ‘Display 1’ in a conference room is a slide projector whose temporal cut-off frequency is close to 0 and ‘Display 2’ in the same room is a plasma display that has a high temporal cut-off frequency, our algorithm will automatically select ‘Display 2’ as the main display for video. Similarly, if ‘Display 1’ has higher resolution than ‘Display 2’, the system would prefer showing a static slide with ‘Display 1’. For other subtle users’ preferences, the system may gradually learn them through online probability updates.

6. EPIC Application Scenarios

EPIC can be used in various ways to enhance a presentation. For long or involved presentations, presenters will often frame the current context with a subtopic slide or an outline slide. If these are presented too often, they may waste presentation time with redundant information. Moreover, users may still lose context if they are inattentive to an outline slide in the main presentation stream. Finally, arranging all slides in one stream is inconvenient for using outline slides to navigate within a presentation.

All these problems can be easily tackled with multiple displays. For example, when two displays are available in a venue, a presenter can put a subtopic slide on a presenter-accessible display and a detailed presentation slide on another large display. The presenter can then provide context to audience members by indicating the current subtopic. The presenter may also navigate the presentation through the subtopic display. Because the subtopic display is always on, the presenter may skip subtopic slides in the main presentation stream.

With multiple displays, supporting images and video may be presented on auxiliary displays. By doing this, a presenter can choose to clarify a text statement on the major display. Because the screen is not taken up with the supporting image, related text will still fit on the slide. The presenter also has more choices to show multimedia data in a short period. In a different scenario, the presenter may consider setting a touch-screen display as an electronic whiteboard, composing surrounding sound for multiple loudspeakers, or dimming room lights.

A benefit of EPIC is that it may be used to schedule device actions in multiple meeting rooms. For example, we can allow EPIC to configure a teleconference environment. At presentation state 0, we may let the software connect camera 1 in room A to display 1 in room B, and connect camera 1 in room B to display 1 in room A. Similarly, we can schedule microphone-speaker connections, camera poses, projector lifts, motorized projection screens, room partition switches, and many other device actions. This kind of configuration need only be configured once for every teleconference environment. With these settings stored in the DST, the software will automatically set up all devices when the presentation is started.

EPIC may also be used for managing conference presentations. By checking that submitted presentations satisfy visual distortion criteria given a particular meeting room arrangement, the system can suggest font and slide changes to make the presentation more visible and thus effective.

7. Results from a Specific Application

Many issues may affect users’ preferences of an authoring result. These issues include the presentation’s content, the user’s presentation style, the meeting environment, and other factors. Given the wide variety of presentation venues and styles, it is not practical to obtain reliable statistics for a general evaluation. However, we report some results we collected from a specific usage scenario.

The meeting environment in our example is a corporate conference room whose top view and setup are shown in Figure 9. The height of the conference room is 110 inches. A corner of the conference room is selected as the origin of our coordinate system. The z axis direction of the coordinate system is from the ground to the ceiling. There are three

displays in the meeting room. Their ground-to-center heights are all 70 inches. Their refresh rates are 75 Hertz. The dimensions and installation heights of these displays are shown in Table 1. Based on display dimensions, resolutions, and simple geometry, it is straightforward to obtain the display parameters shown in Table 2.




Figure 9. Top view of our conference room

Table 1. Display Dimensions and Resolutions

	DISPLAY1	DISPLAY2	DISPLAY3
Height (inches)	72	25	25
Width (inches)	96	44	44
Resolution	1024x768	800x600	800x600

Table 2. Display Parameters

	DISPLAY1	DISPLAY2	DISPLAY3
x_1 (inches)	96~192	26~70	218~262
y_1 (inches)	0	0	0
z_1 (inches)	34~106	57.5~82.5	57.5~82.5
ϕ_1 (direction)	(1,0,0)	(1,0,0)	(1,0,0)
θ_1 (direction)	(0,0,-1)	(0,0,-1)	(0,0,-1)
r_{h1} (pixels/inch)	10.667	18.18	18.18
r_{v1} (pixels/inch)	10.667	24	24

With these parameters, we may compute the perceptual scaling factors for a horizontal line or a vertical line when we know the observer's eye location. According to [11], the average eye height of a person is 46.1 inches. According to [10], the pixel size of a human's fovea may cover 0.31' spatial angle. That is equivalent to $\omega_s = 96 \text{ cycles/deg } ree$. With these data, we can compute α and β variations corresponding to various display portions. For an audience member sitting at location 1 (Figure 9) with average eye height, α and β variations are shown in Figure 10. From Figure 10, it is obvious that the visual distortion is more likely to be introduced in the horizontal direction for people sitting at location 1.

Figure 11 shows three typical slides used in a presentation. We numbered these slides with '1', '2', '3'. We assume the highest slide resolution is 1280 by 960, and all probability distributions are uniform. The computed distortions of various slide arrangements for the audience member at location 1 are shown in Figure 12. Based on these data, it is straightforward to automatically determine the minimum distortion arrangement. The real slide arrangement in the conference room is illustrated in Figure 13.




Figure 10. Scaling factor variations on three displays for a person sitting at location 1. The units on all axes are inches. (a) The variation of α (b) The variation of β .

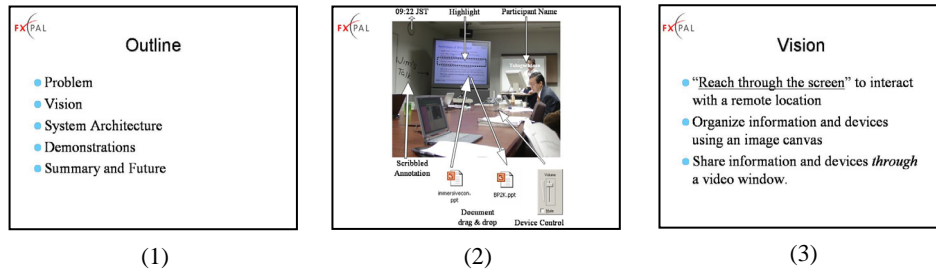



Figure 11. Three typical slides used in a presentation.



Slide Order on Displays '1', '2', and '3'.

Figure 12. Visual distortions caused by various slide arrangements



Figure 13. A typical minimum distortion slide arrangement for a person sitting at location 1.

Since the conference room has three displays available, we start by showing a subtopic or table of contents h-slide on a side display. Our system can automatically detect subtopic h-slides with titles, such as 'subtopic', 'contents', or 'outline' etc. When a subtopic h-slide is not available, the system can automatically generate one based on the slide titles from the PPT file. If an h-slide is an untitled media file, the software will use 'media support' to describe it on the subtopic h-slide. This description could be edited by the user if desired. To automatically generate the presentation, we start by initializing the system as follows:

$$p(\text{on display2} | \text{subtopic h-slide}) = 1$$

$$p(\text{on display3} | \text{prev display} = \text{display1}) = 1$$

All other probability functions are initialized as uniform distributions. The effect of this initialization is to automatically put a subtopic h-slide on display2, put the current h-slide on display1, and put the previous h-slide on display3.

EPIC is now installed on a 2.66GHz/504MB P4 machine in the conference room. It takes about 2 seconds to generate an EPIC presentation from a ~40-slide PPT file. (This delay is mostly due to the PowerPoint API necessary for extracting h-slide thumbnails.) We find this time delay is acceptable to various presenters and audience members.

Since many users are not very familiar with functions provided by EPIC now, most of them let the system automatically generate multi-display presentations. The data collection for EPIC study is still ongoing, and we will present additional data and conclusions in our future paper.

8. Conclusions and Future Work

We have presented a media tool that supports distributed devices. The core contributions are the authoring interface and the media arrangement methods that support authoring from untended automatic to fully manual. With this tool, a presenter can easily use multiple devices in a presentation environment through a simple visual interface. With the help of live video, EPIC also supports remote presentations to a distant venue.

9. References

1. Chunyuan Liao, Qiong Liu, Don Kimber, Patrick Chiu, Jonathan Foote, Lynn Wilcox. Shared Interactive Video for Teleconferencing. *Proc. of ACM MM'03*, pp 546-554.
2. ConferenceXP: <http://www.conferencexp.net/community/>. [10/01/03]
3. Freelance: <http://www.lotus.com/products/smrtsuite.nsf/wPages/freelance>. [10/01/03]
4. Johanson, B., Hutchins, G., Winograd, T. & Stone, M. (2002), PointRight: Experience with flexible input redirection in interactive workspaces. *Proc. of UIST '02, ACM Press*, pp. 227-234.
5. OpenOffice: <http://www.openoffice.org>.
6. Patrick Chiu, Qiong Liu, John Boreczky, Jonathan Foote, Tohru Fuse, Don Kimber, Surapong Lertsithichai, Chunyuan Liao. Manipulating and Annotating Slides in a Multi-Display Environment. *Proc. of INTERACT '03*, pp. 583-590.
7. Richard Anderson, Jay Beavers, Tammy VanDeGrift, Fred Videon, Videoconferencing and Presentation Support for Synchronous Distance Learning. *33rd ASEE/IEEE Frontiers in Education Conference*.
8. Richard Anderson, Ruth Anderson, Tammy VanDeGrift, Steven Wolfman, Ken Yasuhara. Promoting Interaction in Large Classes with Computer Mediated Feedback. *Proc. Of CSCL'03*.
9. Hangjin Zhang, Qiong Liu, Surapong Lertsithichai, Chunyuan Liao, Don Kimber. A Presentation Authoring Tool for Media Devices Distributed Environments. Accepted by ICME2004.
10. Fulton, James T., PROCESSES IN BIOLOGICAL VISION. Available on the Internet: <URL:http://www.4colorvision.com>
11. Ergonomics Data & Mounting Heights. Available on the Internet: URL: http://www.ergotron.com/5_support/literature/PDF/ergonomics_arms_data.pdf
12. Qiong Liu, Don Kimber, Jonathan Foote, Lynn Wilcox, and John Boreczky. FLYSPEC: A Multi-User Video Camera System with Hybrid Human and Automatic Control. *Proc. of ACM Multimedia 2002*, pp. 484 – 492, December 1, 2002.
13. Beatriz Colomina. Enclosed by Images: The Eameses' Multimedia Architecture. Grey Room 02, Winter 2001, pp. 6-29.
14. Library of Congress. The Work of Charles & Ray Eames. URL: <http://www.loc.gov/exhibits/eames/>