Authoring 3D Hypermedia for Wearable Augmented and Virtual Reality

Sinem Güven   Steven Feiner

Department of Computer Science
Columbia University
New York, NY 10027

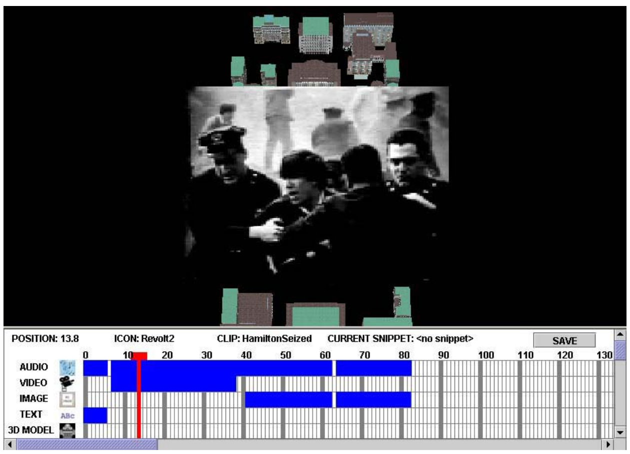

Figure 1. A clip being previewed on the desktop via the presentation component.

Figure 2. A clip being presented in AR by the presentation component.

Abstract

Most existing authoring systems for wearable augmented and virtual reality experiences concentrate on creating separate media objects and embedding them within the user's surroundings. In contrast, designing narrative multimedia experiences for such environments is still largely a tedious manual task. We present an authoring tool for creating and editing 3D hypermedia narratives that are interwoven with a wearable computer user's surrounding environment. Our system is designed for use by authors who are not programmers, and allows them to preview their results on a desktop workstation, as well as with an augmented or virtual reality system.

1. Introduction

Augmented Reality (AR) enhances the user's perception of her surroundings by combining the physical world with a complementary virtual world. AR thus provides an especially rich medium for experimenting with location-aware and context-aware applications, in which virtual objects are designed to take the environment into account.

A number of researchers have explored ways of allowing users to attach individual media objects to the physical world, for display on head-worn displays (e.g., [6, 23, 27]). However, there is currently little support for authoring 3D narrative presentations for AR and virtual reality (VR). Consider, for example, our work on *situated documentaries* [13], which present the mobile user with narrated multimedia stories that are embedded through AR within the same physical environment that the documentaries describe.

Our original situated documentaries were prepared manually by students in Columbia's Graduate School of Journalism, working with students in our lab. Because there was no authoring system for content creators, they had to work in collaboration with programmers who coded the documentaries, and needed to specify in advance which material should be used, how it should be interconnected, and how it should be placed within the surrounding environment. This not only restricted the times at which content creators could work, but stifled their ability to experiment with different approaches because of the programmer effort and turnaround time required. Furthermore, there was little flexibility once the documentary was created, since programmers had to assist with even the smallest changes, making editing and extensions impractical.

To address these issues, we are developing the MARS (Mobile Augmented Reality System) Authoring Tool, which allows content creators to design and edit, without programmer assistance, 3D narrative multimedia experi-

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC'03)
1530-0811/03 $ 17.00 © 2003 IEEE

IEEE COMPUTER SOCIETY

ences that are interwoven with the mobile user's surrounding environment. Desktop authors select media objects and position them spatially, using a 3D model of the world, and temporally, using a timeline, as shown in Figure 1. At any point in its development, the documentary can be viewed in AR or VR, as shown in Figure 2.

In the remainder of this paper, we first describe related research in Section 2. Next, we introduce the MARS Authoring Tool user interface (UI) in Section 3 and its system architecture in Section 4. In Section 5, we discuss our collaboration with colleagues from the Graduate School of Journalism in formative evaluation of the system and the design decisions we made as a result. Finally, we present our conclusions and future work in Section 6.

2. Related Work

There are many research groups that are investigating ways of authoring content for mobile users. The simplest approaches allow users to attach virtual "post-it" notes to physical objects and locations for head-tracked [6, 27] or hand-held displays [20]. Jebara et al. [14] describe how AR and computer vision can be used to associate first-person video and audio recordings with snapshots of real objects that the user sees. A recording is played back when the vision system recognizes the object with which it is associated. Rekimoto et al. [23] present a system that allows users to attach voice notes, or photographs to their surroundings. Fisher [8, 9] has developed an authoring tool for creating outdoor AR experiences, in which authoring is performed on a desktop computer using a web-based 2D map or in the field with a mobile phone. Finally, Rozier et al. [24] describe an audio AR system in which audio recordings can be positioned in the outdoor environment, linked to related ones, and heard as the mobile user encounters them.

While these researchers address ways of associating individual media objects with the user's surroundings, we instead concentrate on authoring documentaries that tie together many different media objects into coherent, interlinked, hypermedia presentations that are embedded within the physical world. A key challenge is that the media objects in a documentary are often interdependent; for example, one object may be related to other objects at the same place, and different objects at other places.

Other researchers have concentrated on creating AR content in 3D. Poupyrev et al. [21] have developed Tiles, an AR authoring system whose users can add, remove, copy, duplicate and annotate virtual objects in the 3D workspace. Sinclair et al. [25] present a tangible hypermedia system using AR. Instead of conventional linking mechanisms, they make use of real world metaphors to author linked 3D hypermedia. For example, users can "sprinkle" on labels or links by picking up a "spice pile marker" and shaking it over regular object markers that are to be labeled or linked. Haringer et al. [10] describe PowerSpace, an AR authoring system, which uses Microsoft PowerPoint to create initial content. The PowerSpace editor allows users to extend this content into 3D presentations and view them in AR. Dörner et al. [5] sketch a high-level approach to authoring AR that exploits reusable building blocks. None of these systems provides support for authoring multimedia documentaries, or embedding them into geographical locations.

Commercial 3D authoring tools, such as Macromedia Director [17], provide a timeline on which to arrange multimedia objects, drag-and-drop manipulation for easy content editing, sophisticated animation mechanisms, and the ability to import/export content from other widely used applications. However, it is not possible to author situated documentaries with any of these systems because they do not make it possible to associate objects or data with geographical locations. Furthermore, these tools require scripting support to make use of their powerful features, such as creating 3D overlays or backdrops. This makes them unsuitable for our purposes, because our aim is to provide an authoring tool for non-programmer content creators. Our work also differs from 3D authoring tools such as 3DM [1], ISAAC [18], RealSoft 3D [22], Corel Bryce5 [3] and Lightwave 3D [15] in that they emphasize building 3D models, as opposed to authoring documentaries or associating information with the user's surroundings.

MacIntyre et al. [16] apply media theory to the development of narrative AR experiences. They create multimedia narratives in which the user can either watch the narrative, or become part of it by interacting with its virtual characters. Currently, they are exploring the possibility of extending existing authoring tools, such as Macromedia Director, to provide support for authoring AR content. However, their authoring approach does not visually support the hypermedia linking and 3D editing we need to create our documentaries. Finally, Holm et al. [11] present a desktop authoring tool for a 3D VR system. They emphasize the development of instructional simulations with rich 3D geometry, controlled by hierarchical event networks created in a visual programming editor, as opposed to authoring and experiencing multimedia documentaries.

3. User Interface

Our MARS Authoring Tool consists of an *authoring component* and a *presentation component*. The authoring component is used to create situated documentaries on a desktop system. The presentation component allows users to explore these documentaries in AR or VR, and preview them on a desktop system. In the rest of this section, we first introduce the building blocks from which situated documentaries are constructed, and show how the presen-

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC'03)
1530-0811/03 $ 17.00 © 2003 IEEE

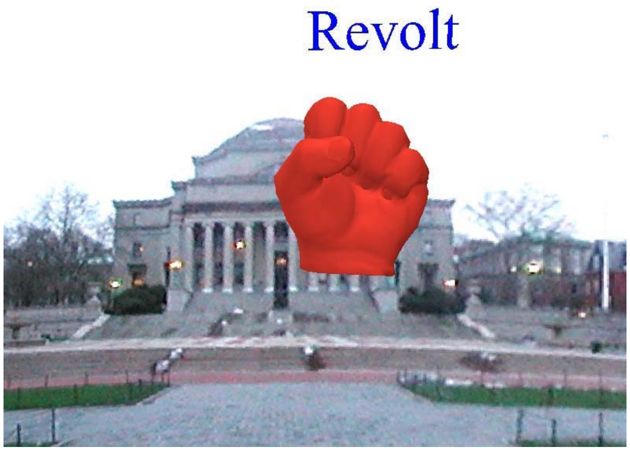

Figure 3. Initial view of a documentary in AR.

Figure 4. Selected icon is highlighted and presents user with contents. Arrow next to clip provides access to linked clips.

tation component UI allows the user to experience documentaries. Then, we present the authoring component UI, and describe how it is used to create and edit situated documentaries.

3.1. Documentary Structure

At the lowest level, a situated documentary created by the MARS Authoring Tool is composed of media objects: audio, video, images, text and 3D models. To be included in a documentary, a media object must be encapsulated as a *snippet,* which contains additional information, such as its start and end times, position, and other data that we describe later. A snippet can be included directly in the 3D environment or it may be added to a named *clip*, which is a collection of snippets. One or more clips are associated with a named *icon*, which is an ordered list of clips, represented visually by a 3D model positioned within the 3D environment. At the highest level, a situated documentary is composed of its icons. Hypertextual links, represented by arrows, allow a clip in any icon to reference one or more clips in any other icons.

3.2. Presentation Component

3.2.1. Application Scenario

The situated documentary shown in this paper was created with our MARS Authoring Tool, from media objects originally developed for an older, handcrafted, situated documentary about the Columbia Student Revolt of 1968 [13]. Our stories are designed to be presented at the actual site using our wearable MARS backpack system [7]. We can also test our system indoors by simulating the site using images captured with a 360º omnidirectional camera [19] and mapped to a hemisphere surrounding the user, as shown in Figure 3. Here, our user virtually stands in the middle of Columbia's College Walk. She can look around to view the documentary's icons, such as the red fist labeled "Revolt," located in front of Low Library, which the author chose to symbolize the strike. The user selects the icon with a wireless handheld trackball, causing it to highlight and display a list of its clips, as shown in Figure 4.

The user then decides to view the *StudentsProtest* clip and selects its name, causing it to highlight and start to play, as shown in Figure 5. Depending on their properties, the snippets that make up the clip will subsequently appear on the head-worn display as screen-stabilized (i.e., fixed within the head-worn screen coordinate system), as shown in Figure 5, or world-stabilized (i.e., fixed within the world coordinate system), as shown in Figure 6.

Figure 5. Screen-stabilized image snippet is fixed at its screen position as user looks around.

Figure 6. World-stabilized image snippet is fixed at its world location as user looks around.

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC'03)
1530-0811/03 $ 17.00 © 2003 IEEE

Figure 7. Screen-stabilized world in miniature indicates user's current location and the Revolt2 icon containing the linked clip. The user has turned toward the Revolt2 icon, but hasn't selected it, so the clip menu for Revolt is still displayed.

A documentary's clips may be distributed across one or more icons. The *Student Revolt* documentary being explored here describes events that started in front of Low Library and then continued at Hamilton Hall. As shown in Figure 4, the arrow to the right of the *MarchOnCampus* clip represents a link indicating that the story continues elsewhere.

In this case, the user decides to follow the link, and activates it by selecting the arrow with the wireless handheld trackball. A screen-stabilized World In Miniature (WIM) [28] appears, showing the user's current location and the 3D icon(s) that contain the linked clip(s). In this case, there is a single linked clip. Figure 7 shows the WIM and the icon associated with the linked clip after the user turns to look at its icon. Next, she selects the icon to view its clips. Figure 2 shows the user viewing the clip following the linked clip. The linked clip appears with an arrow to its left, allowing the user to find the clips that link to it. Figure 1 shows the same clip being played in desktop preview mode. We decided to use a WIM as a part of the presentation UI to give the user an overview of existing icons and their linked clip relationships. We also experimented with paths to guide the user towards icons of interest. However, without a bird's-eye view of the environment, it can be difficult to point at icons, which are not always visible from the user's current location.

In addition to the Student Revolt situated documentary, we also worked with our colleagues in Journalism to create a documentary about the life of Edwin Armstrong, the inventor of FM radio, situated in part in the Columbia building that once housed his lab. Because Armstrong's lab no longer exists, the user is instead immersed in a 3D model of the lab as it once was. In this scenario, the user can interact with three of Armstrong's inventions using the wireless handheld trackball to view their associated clips. Figure 8 shows a frame from one of these clips being played in VR mode (with the icon contents menu hidden) after the user has selected the FM Radio on the left hand side of the table.

Figure 8. A scene from the Armstrong documentary in VR.

3.3. Authoring Component

Because we have only a single outdoor MARS backpack, we decided that the authoring component should be able to be used on a desktop or laptop, to allow multiple users to author material simultaneously on their own computers.

The authoring component presents a 3D model of the author's surroundings. Figure 9 shows the system loaded with a 3D model of the Columbia Campus. The screen-stabilized menu on the left allows the author to create 3D icons representing locations of interest and five different types of multimedia snippets (audio, video, image, text, and 3D model), and also includes 3D navigation facilities.

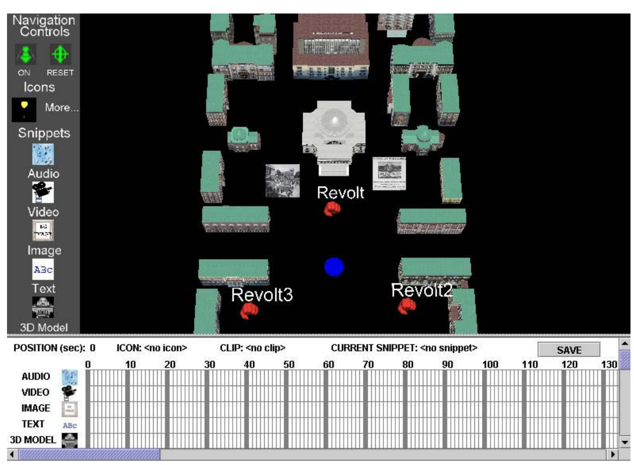

Figure 9. Three 3D icons representing three different locations of interest.

To create our *Student Revolt* situated documentary, the author first identifies the locations at which the events described in the story took place by using the 3D model of the environment. The author then creates three 3D icons (selecting from the collection of 3D icons provided, or choosing her own), and places them in their desired locations.

Next, she names them, *Revolt*, *Revolt2*, and *Revolt3*, as shown in Figure 9. The author then chooses one of the snippet icons on the screen to create a snippet and associates with it the desired multimedia information. The MARS Authoring Tool provides a folder browsing mechanism with thumbnail previewing to help the user choose the correct file. The author can then specify time constraints to indicate when the snippet should start and stop. The end time is automatically calculated for audio and video snippets, whose duration is predefined.

The author can also indicate whether a clip is screen-stabilized or world-stabilized. This important feature makes it possible to create presentations that are quite different from those created with conventional multimedia authoring tools. By default, all snippets are screen-stabilized, which means that during the presentation, regardless of the head orientation of the user, the snippets will appear relative at the same screen location and, thus, will always be visible to the user.

The author can also create world-stabilized snippets by adjusting the position of each snippet, so that it appears at the desired location when it is presented. In our current implementation, world-stabilized snippets are "billboarded" (i.e., texture-mapped onto a 3D polygon that always faces the user).

We decided to use a 3D model of the environment, instead of a 2D map, because our documentaries are intended to be part of the 3D environment. For example, Figures 10 and 11 show an image snippet near the top floor of Low Library, where the events it depicts took place. Using a 3D model makes it easy to associate information with arbitrary parts of the 3D environment. This is an important feature that differentiates our Authoring Tool from 2D map-based hypermedia authoring tools.

Finally, the author previews the created snippet and makes a decision about whether to associate this snippet with the *Student Revolt* documentary or place it by itself in the world. Figures 9–12 show several standalone image snippets embedded in the world so that users can interact with them as they come across them.

More typically, a snippet can be associated with the currently open clip, opening a clip if necessary by interacting with its icon or creating a new clip. The author then drags and drops the snippet onto the clip's icon. When a snippet is added to a clip, it appears on a 2D screen-stabilized timeline that visually represents the clip's contents, as shown in Figure 1. Each snippet's location is determined by its multimedia type and temporal properties. The author can interact with the timeline to move the snippets horizontally to change their temporal properties, delete snippets, or edit their contents. As in desktop animation editors, the timeline uses a moving vertical cursor to indicate the current position in a clip being played.

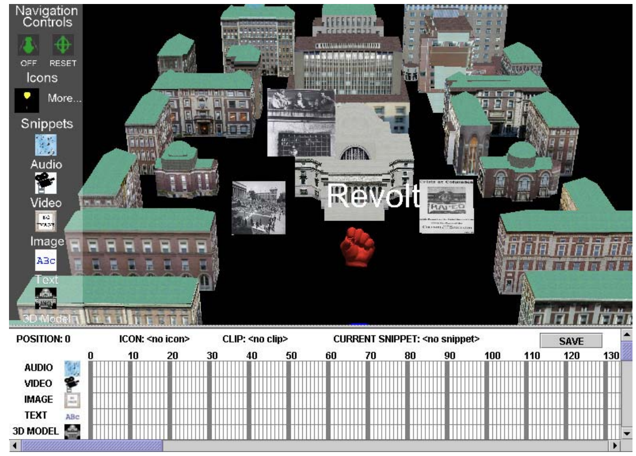

Figure 10. Image snippet associated with the top floor of Low Library.

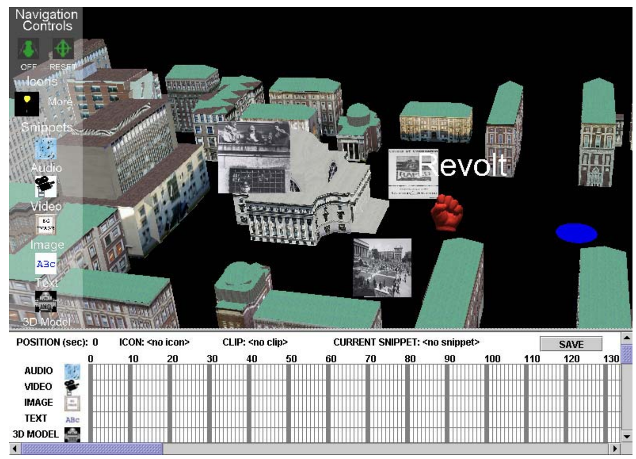

Figure 11. The author can zoom in and navigate within the 3D environment of the MARS Authoring Tool for precise positioning of snippets. Snippets and 3D icons are billboarded to always face the user.

After the author creates snippets and adds them to a clip, the clip can be saved, associating it with the 3D icon that contains it, making it part of its situated documentary. When a clip is associated with an icon, an XML file is created to hold information specific to the clip, as described in Section 4.2.1. The author can then close the clip and view it in the presentation component.

3.3.1. Creating and Linking Sub-Documentaries

Part of what makes our documentaries interesting is that they can use hypertext links (e.g., to connect events that took place at different locations). For example, in our

1530-0811/03 $ 17.00 © 2003 IEEE

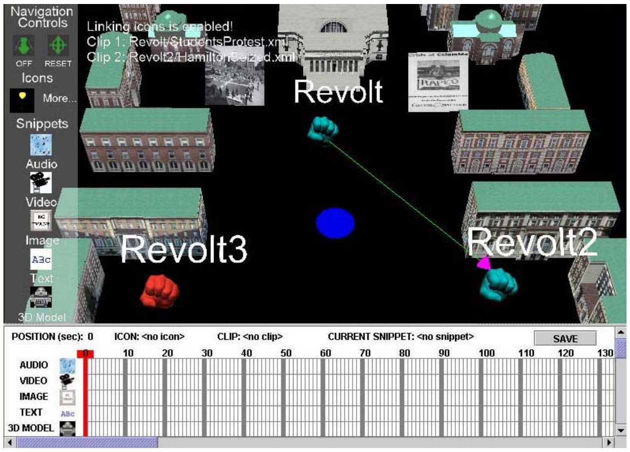

Figure 12. Linking clips associated with icons representing locations in the authoring component.

original work on hand-coded situated documentaries [13], icons (depicted as flags on poles) were linked together into a story through *nextFlag* relationships. This made it easy for a user viewing the portion of a documentary associated with a flag to select "next flag" from a menu to be directed to the flag at which the next part of the story resides. While the user was not obligated to explore that flag next, even if they selected the menu item (i.e., the flag's contents were not automatically invoked), this approach did require that authors specify a single preferred path through their documentary. The relationship was also between icons (flags), rather than between individual clips associated with the icons.

To provide greater flexibility, we allow the author to specify *nextClip* relationships between clips, with any number of these relationships allowed from or to any clip. A nextClip relationship is created by selecting the source and destination clips and is depicted by a 3D arrow between the clips' icons, as shown in Figure 12. To avoid confusion when specifying these relationships, the only arrows depicted are those associated with the clips being connected.

4. System Design and Architecture

4.1. Hardware

We are developing our system on a 1.67GHz AMD Athlon MP 2000+ PC with an NVIDIA Quadro4 750 XGL graphics card and 1GB RAM, running Windows XP. We use a Sony LDI-D100B stereo, color, SVGA, see-through head-worn display with adjustable transparency, tracked by a 6DOF Intersense IS-900 wide-area tracker for indoor, wearable AR and VR. When viewing situated documentaries outdoors using our MARS backpack [7], we use an Ashtech GG24 Surveyor real-time kinematic GPS system for position tracking and an

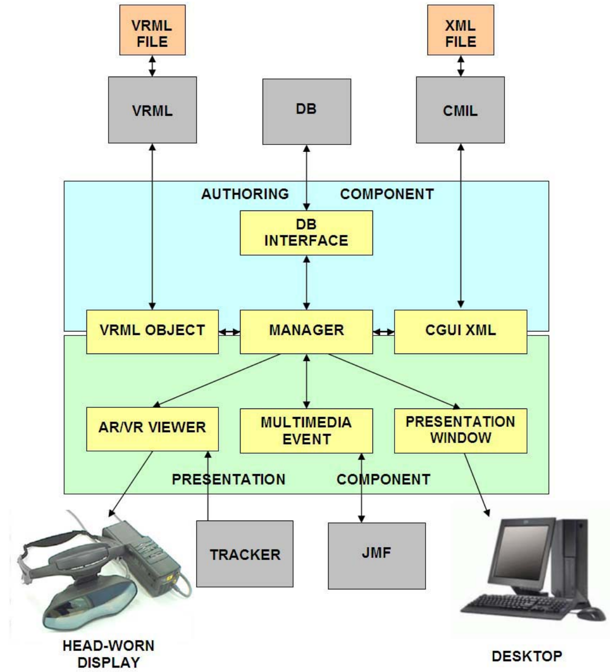

Figure 13. System Architecture. Main classes are shown within authoring and presentation components.

Intersense IS-300 Pro hybrid orientation tracker on the head-worn display.

4.2. Software Infrastructure

The MARS Authoring Tool is written in Java and Java3D, together with SQL DB support for storing information about documentaries and their contents, and XML support for encoding documentaries and their relationships with the world. All 3D icons and WIM contents are VRML 97 [2] models loaded into Java3D. We chose VRML 97 not only because it is a formal standard for 3D representation, but also because its human readable format makes simple adjustments easy and there are many freely available modeling tools. The playback mechanism is implemented using the Java Media Framework (JMF), while our 2D interaction mechanisms, such as the timeline, use Java2D. Figure 13 provides a system overview. The Manager component sets up the scene graph and provides a means of communication between system components. The remaining system components are described in the following sections.

4.2.1. XML Component

CMIL (Contextual Media Integration Language) [4] is an XML-derived language used to describe the relationships between digital multimedia files and contextual information. When we designed the XML component of the system, CMIL and SMIL [12] were the two top candi-

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC'03)
1530-0811/03 $ 17.00 © 2003 IEEE

```xml
<?xml version="1.0" encoding="UTF-8" ?>
<cmil++>
  <head>
    <title>Sample Clip</title>
  </head>
  <body>
    <group>
      <node id="1">
        <attributes>
          <time begin="0.0" end="3.9" />
        </attributes>
        <media>
          <slide src="Image1.jpg" />
        </media>
        <effect>
          <behavior ="flip" />
        </effect>
        <location>
          <position x="140" y="162" z="20"
           type="world" />
        </location>
      </node>
    </group>
  </body>
</cmil++>
```

Figure 14. Sample CMIL++ file.

dates. We chose CMIL because it had better DTD support for our work. Nevertheless, we needed to extend it, as described below, to create what we call CMIL++.

A CMIL++ file represents the multimedia snippets that comprise a clip. Figure 14 shows a CMIL++ file that contains a single image snippet. A CMIL++ file is named after its clip, and each snippet that it contains has the following properties:

- a unique snippet ID.
- attributes, such as the snippet's start and end times.
- information about the type of media the snippet holds (audio, video, image, text, 3D model) and the name of the media file.
- behavior information, specifying how the snippet will appear and disappear from the screen.
- information about the screen position at which the snippet should appear, if it is screen-stabilized.
- location information, describing the position at which the snippet is located, and whether it is screen-stabilized or world-stabilized.

The original CMIL implementation did not support transition effects, such as dimming or flipping an image snippet when it ends. CMIL also did not support specifying a snippet's location, along with whether the snippet is screen-stabilized or world-stabilized. Therefore, we added <effect> and <location> nodes. We then wrote the *CGUI XML* interface (Figure 13) to CMIL++, which is used to store all created clip information in the appropriate CMIL++ file. The *CGUI XML* interface also makes it possible to edit a CMIL++ file while the relevant clip is being edited.

4.2.2. JMF Component

The Java Media Framework (JMF) API enables audio, video and other time-based media to be added to Java applications and applets. This makes it possible to capture, play back, stream and transcode multiple media formats.

We wrote *Multimedia Event* (Figure 13) as a playback interface to JMF, and use it to view the entire documentaries or individual snippets before or after associating them with a documentary. The JMF API also computes duration information for multimedia attached to audio and video snippets, which is important for authors when composing documentaries.

4.2.3. DB Component

We use Microsoft SQL Server to manage our database, which contains all details of a documentary. An *Icon* table stores information about icons and their physical locations. The *Multimedia* table stores information about snippets and their relationships to clips. Information about clips and their associations with icons is stored in the *Presentation* table. Finally, the *Links* table maintains the *nextClip* relationships.

5. Formative Evaluation

During the early phases of this project, we collaborated with colleagues from the Graduate School of Journalism at Columbia University to evaluate the usability and performance of the MARS Authoring Tool. Eleven Journalism students used the first version of our tool to create situated documentaries as part of their Media Studies coursework. The students gave us feedback on what they would have liked beyond the functionality we had provided, what they found helpful, and what they found difficult to use. We identified and addressed some important UI and usability issues based on their feedback.

The students confirmed that the MARS Authoring Tool eliminated the need for a programmer to turn their ideas into an implementation, and, therefore, that they could spend more time designing documentaries, rather than helping to implement them. The MARS Authoring Tool also gave them the flexibility to experiment with different ideas. They especially liked being able to create documentaries whenever and wherever they liked, and being able to modify or extend them easily.

One of the problems that Journalism students reported was that the system was not as flexible as they wanted it to be for content creation. They wanted to include complex animations that our system did not support, and, therefore, felt restricted during content creation. The students were accustomed to using commercial multimedia authoring tools, such as Macromedia Director [17], in their other courses (and professional careers), instead of

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC'03)
1530-0811/03 $ 17.00 © 2003 IEEE

the restricted functionality offered by our authoring tool. Yet the existing tools did not support the creation of situated documentaries.

Rather than trying to compete with commercial tools, we decided to encourage the students to use them to create snippets, by exporting a video file that becomes the snippet's contents. This makes it possible to take advantage of the rich animation facilities that these tools provide. However, it is not possible to edit the imported video within the MARS Authoring Tool.

The Journalism students also reported being confused by the presentation component's hierarchical menu system used in our earlier prototype and by the dialogue box we used to supply timing information for a snippet to position it temporally within a clip. We addressed these issues by using a single-level screen-stabilized menu and implementing the interactive timeline editor described previously.

6. Conclusions and Future Work

We have presented an authoring system for creating and editing situated documentaries that are spatially interwoven with the real world and interconnected by hypertextual links. These 3D hypermedia presentations are experienced by wearable computer users as they explore their environment. Our MARS Authoring Tool supports novice users, eliminating the need for programmers to laboriously turn authors' ideas and raw media material into running presentations, as was required in our original work on situated documentaries.

There are many directions that we would like to explore in future work. To avoid the potential authoring bottleneck posed by our single outdoor MARS backpack, we designed our authoring component to run on conventional desktop and laptop computers. We are currently extending the UI to provide support for authoring situated documentaries directly on-site in AR and VR using a head-worn and a hand-held display in combination.

Now that we have a relatively straightforward way for content experts to turn their ideas into presentations, we would like to address some of the problems that arise when the world is populated by more than the handful of presentations previously created in our collaborative efforts. For example, we intend to filter the presentations shown to the user, based on factors such as their current location and interests. We are also interested in making use of knowledge-based mechanisms to determine which of a clip's set of linked clips to show (e.g., based on the user's location, interests, and previous interactions).

We would also like to explore the potential for making it possible for snippets, clips, and icons to change position and scale over time, and to react in more general ways to user interaction. Our original handcrafted situated documentaries supported the simultaneous use of two different displays (head-worn and hand-held) by a user interacting with a documentary, and included Java applets (e.g., to present a graphical timeline on the hand-held display that faded up and down 3D models of surrounding buildings on the head-worn display, based on the year selected) [13]. Making these kinds of capabilities available to users with little or no programming ability will be challenging.

Finally, our use of XML suggests many possibilities for 3D hypermedia authoring, such as treating the world itself as the scaffolding for a physically situated World Wide Web [26] to be explored in situ.

7. Acknowledgments

Students in Prof. John Pavlik's courses in Columbia's Graduate School of Journalism created material shown in the situated documentary images. This research was funded in part by Office of Naval Research Contracts N00014-99-1-0394, N00014-99-1-0683, and N00014-99-1-0249, NSF Grant IIS-00-82961 and IIS-01-21239, and a gift from Microsoft.

8. References

[1] J. Butterworth, A. Davidson, S. Hench, and T. M. Olano, "3DM: A Three Dimensional Modeler Using a Head-Mounted Display," *In Proc. Symp. on Interactive 3D Graphics*, Cambridge, MA, March 1992, 135-138.

[2] R. Carey and G. Bell, "The Annotated VRML 97 Reference", Addison Wesley Longman, Inc., 1997–1999.

[3] Corel Bryce 5, Corel Corporation, Ontario, Canada, www.corel.com.

[4] R. Dietz, "CMIL Specification 0.9," Oacea, Inc., Bloomington, IN, www.oacea.com/cmil.

[5] R. Dörner, C. Geiger, M. Haller, and V. Paelke, "Authoring Mixed Reality—A Component and Framework-Based Approach," *In Proc. IWEC '02* (*Int. Workshop on Entertainment Computing*), Makuhari, Japan, May 14–17, 2002, 405–413.

[6] S. Feiner, B. MacIntyre, M. Haupt, and E. Solomon, "Windows on the World: 2D Windows for 3D Augmented Reality," *In Proc. UIST '93* (*ACM Symp. on User Interface Software and Technology*), Atlanta, GA, November 3–5, 1993, 145–155.

[7] S. Feiner, B. MacIntyre, T. Hollerer, and A. Webster, "A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring the Urban Environment," *In Proc. ISWC '97* (*First IEEE Int. Symp. on Wearable Computers)*, Cambridge, MA, October 13–14, 1997, 74–81.

[8] S. Fisher, "Environmental Media: Accessing Virtual Representations of Real-Time Sensor Data and Site-specific Annotations Embedded in Physical Environments," *In Proc. WSMM '01* (*Seventh Int. Conf. on Virtual Systems and Multimedia*), Berkeley, CA, October 25–27, 2001, 407–418.

[9] S. Fisher, "An Authoring Toolkit for Mixed Reality Experiences," *In Proc. IWEC '02* (*Int. Workshop on Entertainment Computing*), Makuhari, Japan, May 14–17, 2002, 487–494.

[10] M. Haringer and H. T. Regenbrecht, “A Pragmatic Approach to Augmented Reality Authoring,” *In Proc. ISMAR ‘02* (*IEEE and ACM Int. Symp. on Augmented Reality*), Darmstadt, Germany, September 30–October 1, 2002, 237–245.

[11] R. Holm, E. Stauder, R. Wagner, M. Priglinger, and J. Volkert, “A Combined Immersive and Desktop Authoring Tool for Virtual Environments,” *In Proc. VR ‘02* (*IEEE Virtual Reality Conf. 2002*), Orlando, FL, March 24–28, 2002, 93–100.

[12] P. Hoschka, "Synchronized Multimedia Integration Language (SMIL) 1.0 Specification," W3C Recommendation, 15 June 1998.

[13] T. Höllerer, S. Feiner, and J. Pavlik, “Situated Documentaries: Embedding Multimedia Presentations in the Real World,” *In Proc. ISWC ’99 (Third IEEE Int. Symp. on Wearable Computers),* San Francisco, CA, October 18–19, 1999, 79–86.

[14] T. Jebara, B. Schiele, N. Oliver, and A. Pentland, “DyPERS: Dynamic Personal Enhanced Reality System,” *In Proc. 1998 Image Understanding Workshop*, Monterey, CA, November 20–23, 1998, 1043–1048.

[15] Lightwave 3D, NewTek, San Antonio, TX, www.newtek.com.

[16] B. MacIntyre, J. Bolter, E. Moreno, and B. Hannigan, “Augmented Reality as a New Media Experience,” *In Proc. ISAR ’01* (*IEEE and ACM Int. Symp. on Augmented Reality*), New York, NY, October 29–30, 2001, 197–206.

[17] Macromedia Director, Macromedia, Inc., San Francisco, CA, www.macromedia.com/director.

[18] M. Mine, “ISAAC: A Meta-CAD System for Virtual Environments,” *Computer-Aided Design,* 29(8), 1997, 547–553.

[19] S. Nayar, “Catadioptric Omnidirectional Camera,” *In Proc. CVPR ’97 (IEEE Conf. on Computer Vision and Pattern Recognition*), Puerto Rico, June 1997, 482–488.

[20] J. Pascoe, “Adding Generic Contextual Capabilities to Wearable Computers,” *In Proc. ISWC ’98* (*Second IEEE Int. Symp. on Wearable Computers*)*,* Pittsburgh, PA, October 19-20, 1998, 92–99.

[21] I. Poupyrev, D. Tan, M. Billinghurst, H. Kato, H. Regenbrecht, and N. Tetsutani, “Tiles: A Mixed Reality Authoring Interface,” *In Proc. INTERACT ’01 (Eighth IFIP TC.13 Conf. on Human-Computer Interaction)*, Tokyo, Japan, July 9–13, 2001, 334–341.

[22] Realsoft 3D, Realsoft Graphics, Finland, www.realsoft.fi.

[23] J. Rekimoto, Y. Ayatsuka, and K. Hayashi, “Augment-able Reality,” *In Proc. ISWC ’98* (*Second IEEE Int. Symp. on Wearable Computers*)*,* Pittsburgh, PA, October 19-20, 1998, 68–75.

[24] J. Rozier, K. Karahalios, and J. Donath, “Here & There: An Augmented Reality System of Linked Audio,” *In Proc. ICAD ’00* (*Int. Conf. on Auditory Displays*), Atlanta, GA, April 2–5, 2000, 63–67.

[25] P. Sinclair, K. Martinez, D. Millard, and M. Weal, “Links in the Palm of your Hand: Tangible Hypermedia using Augmented Reality,” *In Proc. HT ’02* (*ACM Conf. on Hypertext and Hypermedia*), College Park, MD, June 11–15, 2002, 127–136.

[26] J. Spohrer, “Information in Places,” *IBM Systems Journal,* 38(4), 1999, 602–628.

[27] T. Starner, S. Mann, B. Rhodes, J. Levine, J. Healey, D. Kirsch, R. Picard, and A. Pentland, “Augmented Reality Through Wearable Computing,” *Presence,* 6(4), August 1997, 386–398.

[28] R. Stoakley, M. Conway, and R. Pausch, “Virtual Reality on a WIM: Interactive Worlds in Miniature,” *In Proc. CHI ’95 (ACM Conf. on Computer Human Interaction),* Denver, CO, May 7–11, 1995, 265–272.

Proceedings of the Seventh IEEE International Symposium on Wearable Computers (ISWC’03)
1530-0811/03 $ 17.00 © 2003 IEEE