Narrative Event Adaptation in Virtual Environments

Karl E. Steiner
University of North Texas
General Academic Building, Room 320
Denton, Texas 76203
steiner@cs.unt.edu

Jay Tomkins
University of North Texas
General Academic Building, Room 320
Denton, Texas 76203
jt0043@cs.unt.edu

ABSTRACT
There is a tension between user and author control of narratives in multimedia systems and virtual environments. Reducing the interactivity gives the author more control over when and how users experience key events in a narrative, but may lead to less immersion and engagement. Allowing the user to freely explore the virtual space introduces the risk that important narrative events will never be experienced. One approach to striking a balance between user freedom and author control is adaptation of narrative event presentation (i.e. changing the time, location, or method of presentation of a particular event in order to better communicate with the user). In this paper, we describe the architecture of a system capable of dynamically supporting narrative event adaptation. We also report results from two studies comparing adapted narrative presentation with two other forms of unadapted presentation - events with author selected views (movie), and events with user selected views (traditional VE). An analysis of user performance and feedback offers support for the hypothesis that adaptation can improve comprehension of narrative events in virtual environments while maintaining a sense of user control.

Categories and Subject Descriptors
H.5.1 [**Information Interfaces and Presentation**]: Multimedia Information Systems - *Artificial, augmented, and virtual realities.*

General Terms: Design, Human Factors

Keywords: Intelligent Multimedia, Adaptive Interface, Virtual Environment, Narrative and Storytelling

INTRODUCTION
Narrative is an important method of human communication. We tell stories to communicate, to educate, and to entertain. Narratives allow us to reach diverse audiences, to approach difficult subject matter, and to reach audiences emotionally. As new technologies have been introduced, our methods of storytelling have adapted. Narratives were among the initial and most popular content types as books, radio, movies and television were introduced to the public. Virtual environments (VEs) present one of the newest technologies to be applied to the presentation of narrative.

IUI'04, January 13–16, 2004, Madeira, Funchal, Portugal.

Interactive systems create new opportunities in the presentation of narratives. Unlike written or cinematic narratives where the author controls the pace and manner of exposition, in an interactive narrative, the user may assume control of various aspects of the environment. For example, in a hypertext narrative, a user may choose the order in which certain passages are read by navigating hyperlinks. Similarly, in a narrative presented in a virtual environment, a user may choose the order in which events are viewed by navigating through a 3D space. In more sophisticated interactive environments, the user may be able to interact with characters and objects in the environment, playing the role of a character in the narrative, and influencing the flow of the plot.

However, this flexibility introduces a number of issues. Once the user assumes control of viewpoints and interactions, the sight or even occurrence of certain events may no longer be a given. Important events in the VE may occur out of the sight lines of the user due to distance or occluding objects. Users may miss otherwise visible events due to distractions in the virtual environment. Users may see important events yet fail to recognize their significance due to a lack of context. Or in the case of an interactive environment, events may fail to occur altogether due to the user failing to initiate certain interactions.

Designers of narrative VEs have adopted a number of strategies in order to minimize the impact of these issues. In general, these strategies result in tightly scripted and linear narrative, providing little or no opportunity to influence the plot of the story. Since the environment has no awareness of narrative goals, the designer must take measures to force users into viewing narrative events. This usually means either eliminating interactivity at certain points and forcing users to watch pre-defined cut-scenes, or keeping interactivity, but eliminating choices so that the user must ultimately proceed down a particular path or choose a particular interaction. These solutions are unsatisfying for many users, and time consuming for designers to construct.

A fundamental issue behind the limited interactivity is the tension between the user's desire to freely explore and interact with the environment, and the designer's goals of having users experience particular educational, entertaining or emotional events. This has been referred to by others as the balance of writer flexibility vs. user flexibility [9], and the conflict between story coherence and user control [14]. Advances in visualization technology have exacerbated the issue. Virtual environments are now capable of modeling more objects, behaviors, and relationships than a user can readily view or comprehend. Techniques for automatically selecting what events or information to present and how to present them are not available, so VE designers must explicitly specify the time, location, and conditions under which events may occur. As a result, most VEs support only rather linear narratives or scenarios (if they support them at all).

One way to make narrative VE technology more accessible would be to shift some of the burden of managing the presentation from the designer and the user to the computer. Whereas most virtual environments attempt to achieve their goals by directing user attention towards certain locations or viewpoints, we believe a more effective approach would be to adapt events to fit into the users current context where possible. Rather than manipulating users into moving to certain locations or requiring them to adapt particular viewpoints (even dynamically generated viewpoints), the system should take advantage of the locations, views and activities that already occupy the users' attention.

Such an approach could have a number of benefits. With event presentation management, users could become aware of important narrative events regardless of when and where the events occur. Users would receive this information without compromising their immersion or their feelings of control. The process of designing VE narratives could be made simpler. Authors and designers could define event characteristics, then let the presentation manager handle exceptions and conflicts.

1. RELATED WORK

Our work builds off of related efforts in interactive narrative, particularly those situated in virtual environments [8, 13]. While our current interest focuses on adaptation of narrative events for multimedia VEs, we have gained insights from work done on automated cinematography, adaptive narrative hypermedia, and narrative VEs.

Work on automated cinematography and 3D Camera control has addressed issues related to dynamically managing views. For example, ConstraintCam [2] dynamically generates optimal camera angles to support explicitly stated user viewing or task goals. Other approaches, such as CamDroid [4] and the Virtual Cinematographer [7], also explore approaches for intelligent camera control. While these systems address the need to free users and designers from explicitly specifying views for a given event, they do not address issues of event representation.

Recent work on narrative hypermedia has provided insights into a number of areas. Weal describes using the Fundamental Open Hypertext Model to create context-based linking for adaptive narrative hypermedia [20]. A number of projects have explored the integration of cinematography conventions into hypermedia [10, 19]. Other hypermedia systems have dealt with control or direction of narrative. For example, in Carmen's Bright Ideas [11], a director seeks to maintain story structure coherence and fidelity to pedagogical goals while a cinematographer manages multimedia presentation. While this work shares some similar goals with our own, their target environment involves turn-based dialog with 2D animations rather than the real-time 3D interactions we envision. The PEGASUS system [15] includes a presentation manager that can select images (stage), color and sound to enhance the presentation of textual narratives (mythological encyclopedia entries). Gershon describes how narrative can be used to guide data visualization [6]. Presentations of temporal or spatial data could be self-organizing, arranging and rearranging the data based on different narrative perspectives or emphases

There has also been work on VE environments that seek to model narratives and that include an automated "director" responsible for managing selection and flow of events. In the OZ project [2, 21], a director gives instructions to actors in a VE seeking to maximize the dramatic potential of a situation. In more recent work, the Mimesis project [14, 22] includes a director that works to detect and prevent potential plot conflicts introduced by character interactions. Mimesis also includes its own cinematic camera controller [1]. The Experience Learning System being developed by the Institute for Creative Technologies includes research in immersive VR, training, and emotional engagement of user [12, 18].

A related project with goals similar to our own is Galyean's work on narrative guidance of interactivity [5]. Galyean developed a system that selects events and adapts them in order to meet author defined narrative goals. Events are enabled based on the VE world-state, including the users view. Event presentation is adapted using cinematic techniques such as close-ups, and establishing shots. The work resulted in a proof of concept prototype (Dogmatic). Our system has similarities to Galyean's, particularly in terms of overall goals and in the management of events. Our system differs in that we have elected to leave view management entirely in the users control (which precludes the use of cinematic camera techniques), and our event model supports simultaneous events at multiple locations.

Underlying the development of many of these applications is an assumption that users will respond more positively, be more engaged, and have enhanced experiences as a result of interacting with these VE narrative experiences (as opposed to similar non-interactive experiences). These projects are compelling, and many provide anecdotal support for the efficacy of their approach. The ICT project has reported that user studies to measure the impact of emotion on immersion are in progress [12], but their results are not yet available. So to date, quantitative study of user responses to these systems has been limited, and a number of fundamental questions regarding many aspects of the user experience of VR narratives remain unanswered. As an initial step in the development of our platform for narrative VE, we chose to explore user response to event adaptation in a narrative virtual environment.

2. SYSTEM ARCHITECTURE

We are developing a platform for managing and presenting interactive narratives in a VE [16]. This platform currently includes a 3D engine, VE world state management, story state management, and event adaptation (Figure 1). Work is also underway on dynamic event generation and user modeling components. An earlier version of this platform was used to support our exploration of children's collaboration in narrative environments [17].

As our interest in VE focuses more on interaction and less on 3D rendering, we decided to take advantage of a commercial 3D game engine. We use the A5 3D engine from Conitec. The engine and related development tools support 3D world-building and display, allowing us to devote the majority of our attention to developing support for narrative events, interactions and presentation.

We have developed data structures and sensors that allow us to track key information regarding the conditions in the 3D environment (world state). This includes immediate information such as object and user location, as well as relationships between objects (e.g. Is the user near LOCATION-1? Can the user see CHARACTER-3?). Additional information about the user is managed as well, including a history of activities.

We have also developed capabilities for managing and presenting narrative events (story state). An event describes character movements and character-character or character-object interactions. An event is composed of an ordered set of primitive actions that make up the event, as well as meta-data for further categorizing the event and its presentation. Also associated with each event are a set of conditions (triggers) that determine when the event can be presented. These conditions include time, location, other elements of world or story state, or combinations of the above. Once an event has been activated, the adaptation manager selects an appropriate way in which to present the event.

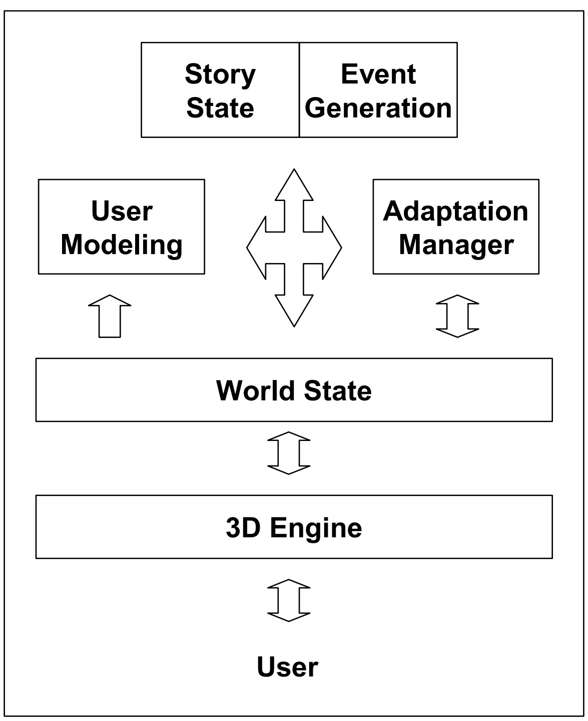

Figure 1. System Architecture

Table 1. Scenario Event Examples

EVENT EXAMPLES		
EVENT	**ACTION**	**ADAPTATION**
A hungry rabbit eats the last of his carrots and goes looking for something else to eat.	Animation of rabbit eating carrots; Sounds of rabbit eating carrots	Amplify sound of rabbit eating carrots (to draw attention to rabbit and convey idea that someone is eating something).
The rabbit notices a snowman's carrot nose and steals it.	Animation and sounds of rabbit stealing carrot nose	Delay actions until user is present at the characters' location (so that user may view and hear this key event).
The snowman follows the rabbit to the lake, but the rabbit has already eaten the carrot. The only item he can find is a fish, so he uses that as a temporary replacement.	Sounds of snowman and rabbit discussing options; Animation of snowman picking up fish and using it as nose	If the user was not present to see the event, the rabbit talks to himself and recounts the event within the hearing distance of the user.

The adaptation manager is a rule-based system that utilizes knowledge regarding the user, the world state, and the event in order to select an appropriate presentation. Some significant user and world-state facts utilized by the system include the user location and orientation, the location and orientation of other characters and props. Significant information regarding events includes constraints on the event such as whether the event must be viewed directly, whether the event must occur at a particular location, and if the event must occur at a specific time. The system also maintains a set of adaptations available for the specific actions that make up an event.

Our current system supports three general categories of event adaptation: time, location, and form (Table 1). A time adaptation can postpone the initiation of an event. For example, the revelation of an important clue might be postponed until the user is present to see it. Location adaptations can change where an event occurs. For example, a conversation between two supporting characters might be relocated to a location closer to the user to enhance the chance that it will be overheard. Our current system supports two general types of changes in the form of an event: audio-visual enhancements and intermediaries. For example, the sound of an explosion, or the size of a plume of smoke might be increased, or a by-stander might approach the user and mention that there was an explosion. The rules for changing the timing and location of events are general-purpose and can be applied to virtually any event. Changes in form are more event-specific, and may require the creation of additional media in order to be fully supported. For example, if an intermediary is to communicate with the user regarding an event, alternative audio or descriptive text must be made available.

Two more components are currently under development: a user modeler and a dynamic narrative generator. The user modeler will generate "belief" and "interest" models of the user, using information such as a user profile, a history of user activities within the VE, and other world state and story state data. Information generated by the user model might include the user's apparent interests within the narrative VE (e.g. the user is seeking treasure and wealth), what the user believes about the narrative or the VE (e.g. the user believes there is a treasure in the hidden room, the user believes that the dog is hostile and will attack), and whether the user falls into a recognized "type" (e.g. the user is an "explorer" rather than a "fighter").

The narrative generator will allow the system to go beyond adapting pre-written events, to generating new events to satisfy storytelling goals. This would allow the system to support plots that adapt to user activities, as well as plots that are randomly generated, providing unique experiences on each viewing.

3. USER STUDIES

A VE with a completely scripted timeline of views and events (similar to a movie) gives the author control over when and how events are presented. It makes intuitive sense that the viewer of such a system would have a high level of comprehension of the significant plot events. However, the viewer would also be likely to feel as if they have little control over the manner in which they view events. Allowing the user to freely explore the environment and to view events as they encounter them would give the user an increased feeling of control and ownership of the experience. However, given a scripted timeline of events, the user runs a substantial risk of missing key plot points by not arriving at the right locations at the right times.

We conducted a pair of studies to explore the efficacy of event adaptation in virtual environments and to uncover some of the issues surrounding the design and development of such a system. We began with a preliminary user study of a smaller pilot system. The feedback we gained from this initial work led us to revise our system. We followed up with another user study using a more advanced version of our event adaptation system.

4. PILOT STUDY

We believed that event adaptation would allow us to strike a balance between authorial direction and user control. If mechanisms were in place to manage the presentation of events (dynamically adjusting the time, location, and method of presentation of events), then users could freely explore without missing key plot points. An adaptive system of this type would ultimately support levels of plot comprehension similar to a system with high authorial direction while still giving the user a heightened perception of control.

We devised a study to test our hypothesis. Using an early version of our narrative VE architecture, we created three separate ways of experiencing a story in VE: a scripted timeline of events and author selected set of views similar to watching a movie (MOVIE), a scripted timeline of events and user selected views similar to a traditional first-person VE (VE-SCRIPT), and user selected views with adaptive event presentation (VE-ADAPT). The VE-ADAPT condition combined the interactivity of the VE-SCRIPT condition with intelligent event adaptation.

Figure 2. Pilot Study Scenario Screenshot

Thirty university students (10 in each condition) participated in the Pilot Study. Subjects were given time to become acquainted with the movement controls for the VE before beginning the narrative. Upon completion of the story, the subjects completed a brief questionnaire that included both objective and subjective questions regarding comprehension, feelings of control, and enjoyment. Afterwards, subjects were debriefed, the goals and methods of the study were explained, and subjects were invited to further explore the scenario using the alternate event and view management schemes.

The narrative used in this study was a fairy tale about a hungry rabbit and two snowmen (Figure 2). The narrative was designed to include multiple adaptation types, multiple characters, events that take place at different locations, and events that occur at the same time. However, not all major plot events were supported by adaptation.

An analysis of the results of the questionnaires offered some support for our initial hypotheses. Overall, users in the MOVIE category had the highest levels of comprehension (based on numbers of comprehension questions answered correctly), followed by VE-ADAPT and VE-SCRIPT (Figure 3). While the overall difference between VE-ADAPT and VE-SCRIPT was not large, when examining comprehension on an event-by-event basis, VE-ADAPT users scored much higher than VE-SCRIPT users for all events where an adaptation was available. Both VE-ADAPT and VE-SCRIPT users reported similar moderate feelings of control over the view, while MOVIE users reported a relatively strong lack of control (Figure 5). When asked about story pace, VE-ADAPT users reported far less comfort with the pace of the story than did the MOVIE and VE-SCRIPT users. Users in all conditions gave positive responses regarding enjoyment of the experience, but VE-ADAPT users were more positive in their responses. We felt that these findings were encouraging and offered at least tentative support for our hypothesis that event adaptation could lead to comprehension levels superior to un-adapted events, while still offering levels of control similar to traditional VEs. While navigation was not under user control for users in the MOVIE condition, VE-SCRIPT and VE-ADAPT users were able to chart their own path through the environment. An observation of the actual usage and a review of plots of the users' navigation paths suggests that VE-ADAPT users spent more time in areas significant to the scenario, and did less wandering to "uninteresting" areas.

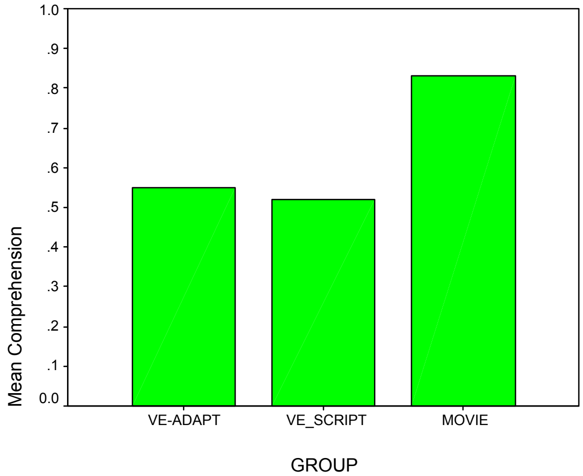

Figure 3. Pilot Study Comprehension

Still, a closer look at the results raised a number of interesting issues. In this implementation, not all events in the VE-ADAPT condition were supported by adaptation. When authoring the story events, adaptations were created for most of the events that were considered interesting, significant, or potentially confusing. Our hypothesis led us to anticipate that VE-ADAPT users would have higher levels of comprehension of events that were supported by adaptation, and this was born out by the data. However, due to time constraints, there were still some important story events that were not supported by adaptation. For events without adaptation, we had expected that VE-ADAPT users would have levels of comprehension similar to VE-SCRIPT users. However, the VE-SCRIPT users had comprehension scores for the un-adapted events that, while lower than the MOVIE users, were higher than those of the VE-ADAPT users. On reflection, we believe that the poor performance on the un-adapted events may be a natural outcome of achieving higher performance on the adapted events. By successfully shifting attention to the adapted events, we may have diverted attention away from the un-adapted events. This suggests that in future implementations, all key events should be supported by adaptation.

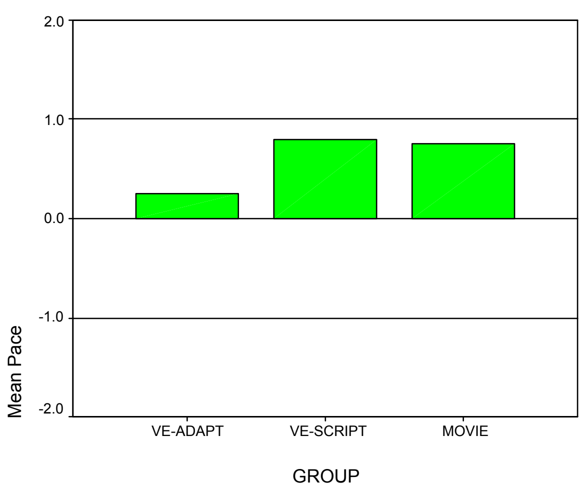

Figure 4. Pilot Study Pace

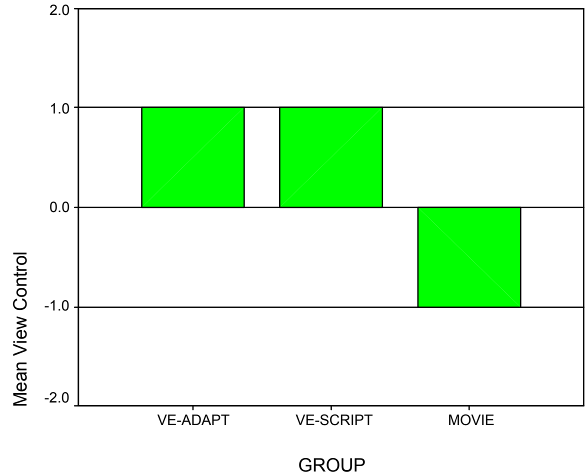

Figure 5. Pilot Study View Control

Another intriguing result relates to the question of story pacing. MOVIE and VE-SCRIPT users responded that they felt the story was better paced than did the VE-ADAPT users (Figure 4). This may be due again to the partial implementation of adaptation. Given the fixed timeline that was present in the VE-SCRIPT and MOVIE versions, the flow of events may have been more obvious. In the Pilot version of the VE-ADAPT scenario, some events were fixed while others were adapted, and this may have resulted in a less obviously satisfying pace for the story. However, the adapted events may have contributed to making the experience more pleasurable. VE-ADAPT users responded most positively to the question of how much they enjoyed the overall experience. Observation and conversations with users during the debriefing support the idea that VE-ADAPT users were engaged with the experience and enjoyed the narrative.

5. USER STUDY

While the results of the Pilot Study were encouraging, they left some open issues. Our limited implementation of adaptation had increased comprehension, but not to the degree to which we had hoped. Users also felt that the story pace in the adapted condition was not as satisfying as the pace in the other conditions. We believed that the limited implementation of adaptation in the Pilot might have had an impact on these results. When a more robust version of our adaptation manager became available, we conducted a follow-up User Study.

5.1 Hypothesis

As in the Pilot Study, we anticipated that users of a VE with adapted events would achieve comprehension levels significantly superior to those experienced by users of a VE with un-adapted events. We also expected adaptation of events would produce a sense of control comparable to that experienced by users of a VE without event adaptation. Further, we expected that using a more robust implementation that provided adaptations for all major narrative events would minimize negative views regarding the pace of the narrative in the adapted condition.

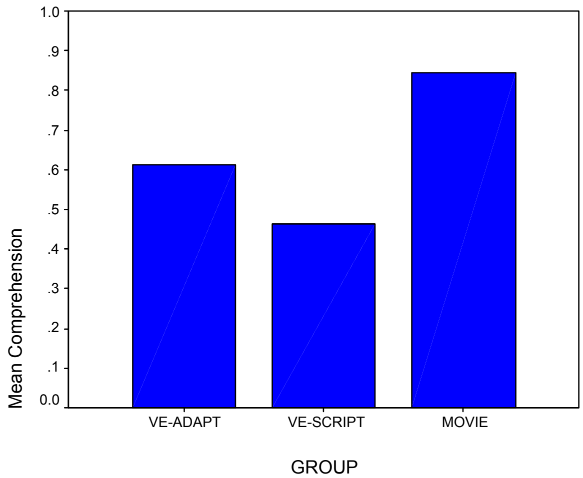

Figure 6. User Study Comprehension

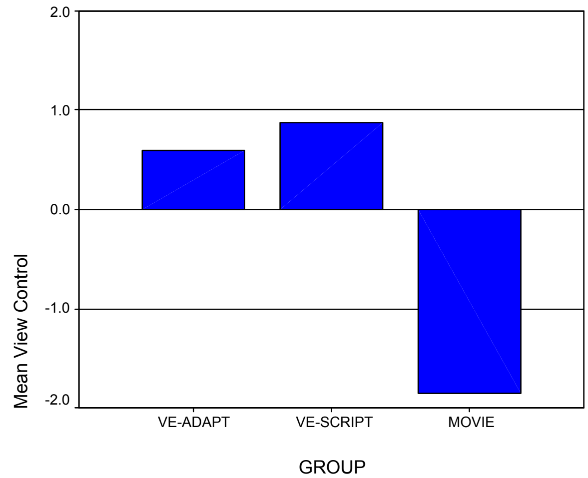

Figure 7. User Study View Control

5.2 Method

Once again, we created a test narrative with multiple characters and simultaneous events at separate locations. Our narrative VE platform was tooled to record user navigation and event timings. We engineered the system to support three separate ways of experiencing the scenario: a scripted timeline of events and author selected set of views (MOVIE), a scripted timeline of events and user selected views (VE-SCRIPT), and an adaptive set of events and user selected views (VE-ADAPT).

Twenty-Eight university students participated in the study. Each was randomly assigned to one of the three categories (10 in VE ADAPT, 10 in VE-SCRIPT, and 8 in MOVIE). They were given a brief introduction to the study, allowed to familiarize themselves with the virtual environment in a practice scenario, then experienced the test narrative with the event and view management dictated by their condition. Upon finishing the scenario, all subjects completed a questionnaire that included both objective and subjective questions regarding comprehension, feelings of control, and enjoyment. Afterwards, subjects were debriefed, and the goals and methods of the study were explained.

5.3 Scenario

The narrative used in this User Study was more complex than the one used in the Pilot Study. More narrative events were supported, dialog was more fully integrated into the storyline, there were more supporting characters, and the settings and props played a more significant role in the narrative. This story related the tale of a clumsy wizard, a greedy king, and a hungry rat. Again, the narrative was designed to include multiple adaptation types, multiple characters, events that take place at different locations, and events that can occur at the same time.

5.4 Results

The results of our analysis of comprehension and feelings of control were consistent with our earlier study, though the differences were more pronounced. Analysis of variance and post-hoc tests revealed significant differences on several measures. The comprehension scores represent the percentage of correct answers to a series of questions based on the narrative events (Figure 6). While the MOVIE condition continued to have higher mean comprehension scores (.84), significant differences ($p < 0.05$) were found between the scores of VE-ADAPT (.61) and VE-SCRIPT (.46).

Users responded to opinion questions by rating their agreement with various statements on a 5-point scale, with –2 meaning Strongly Disagree, and +2 meaning Strongly Agree. Users self-rated their feelings of view control based on their response to the question “I was in control of the way I viewed the story” (Figure 7). Users of both the VE-ADAPT (.60) and VE-SCRIPT conditions (.88) were more positive regarding their feelings of control, and their scores were significantly higher ($p < 0.01$) than those reported by MOVIE users (-1.86). While all three groups reported negative responses to the question “I was in control of the story,” it is interesting to note that the VE-ADAPT users were somewhat less negative in their response (VE-ADAPT: -.70, VE-SCRIPT: -1.29, MOVIE: -2.0). Pacing no longer appeared to be a concern in this implementation, with all users reporting positive levels of agreement with the statement “The story was well paced” (VE-ADAPT: .40, VE-SCRIPT: .43, MOVIE: .71). There were no significant differences on this measure. Similarly, responses to the statement “I enjoyed

Figure 8. User Study Scenario Screenshot

viewing the story" were all modestly positive (VE-ADAPT: .80, VE-SCRIPT: 1.0, MOVIE: .71), again with no significant differences.

5.5 Discussion

With this more robust implementation, we were pleased to see significant results more consistent with our expectations. Testing a larger and more consistently applied set of adaptation with a more complex narrative made the differences in comprehension levels more pronounced. We expected that users in the MOVIE condition would continue to have the highest comprehension scores, but were pleased to see more pronounced comprehension differences between VE-ADAPT and VE-SCRIPT users. We also expected that the MOVIE condition would have the lowest View Control ratings, and we were glad to see that the VE-ADAPT and VE-SCRIPT scores on this measure remained extremely close, indicating that there was little perceived differences in control between these two groups.

Questionnaire comments and interviews with the users following the study confirm these observations. Several VE-SCRIPT subjects commented on their difficulties in following the plot. Some expressed frustration with the sound levels (e.g. they were close enough to characters to hear that a conversation was taking place, but not close enough to hear the exact dialog). Several pointed out that it was challenging to follow the simultaneous activities of multiple characters. A few got engrossed with exploring areas of the VE where no events were occurring, and entirely missed large portions of the plot.

However, while the comprehension levels of VE-ADAPT users were higher than VE-SCRIPT users, our test subjects suggested some areas for improvement. While the VE-ADAPT subjects' response to overall pacing was positive, a few users indicated that event timing could still be improved. For example, one of the adaptations of event timing calls for the Adaptation Manager to delay the event until the user approaches within a specified distance and has a clear line of sight. While in many cases this is sufficient to capture the user's attention, it is possible for a user to run past the location, trigger the adaptation, and continue on without observing the event, and this happened on several occasions. Similarly, the intermediary adaptation called for a minor character to approach the user and deliver dialog explaining an event that occurred out of the user's current sightlines. However, many of the users did not notice the minor character or stop to hear her dialog. Despite the difficulties across the conditions, the general response of users was positive, and several expressed interest in continuing to explore the environment even after the main events of the narrative had concluded.

6. CONCLUSION AND FUTURE WORK

These studies have provided several insights that will guide us in our future development efforts. While the results of the second User Study seem to indicate that event adaptation holds promise as a technique for improving comprehension without sacrificing interaction, successfully balancing user control and author goals remains challenging.

While our current rules for time and location adaptations are flexible and capable of handling most situation, adapting the form of events or communicating them through intermediaries still requires a certain level of hand tuning. We hope to create more general form adaptation capabilities in future generations of the adaptation manager. While agent interaction was not a major focus of this study, we also plan to increase the sophistication of the agents used in our VEs, hopefully allowing them to better command user attention when they have important messages to deliver.

We gained valuable experience in developing the components of our VE architecture and were able to illustrate some of its benefits. For example, our Pilot Study was developed before several major components of our architecture were complete, requiring us to hand-craft events, their triggers, and their adaptations. For our second User Study, more sophisticated Story Management and Adaptation Management components were available. Even though the narrative in the User Study was substantially more complex than the narrative in the Pilot Study, the services provided by the components of our architecture allowed us to develop it in almost half the time (and with superior aesthetic and performance results). Still, while the rule-based adaptation mechanism in the User Study was superior to the ad-hoc adaptation in the Pilot Study (both from the user perspective and in terms of development efficiency), a certain amount of customization of the rules and supporting media was necessary to produce acceptable performance.

We plan to continue our development efforts by exploring ways of allowing event adaptation to automatically deal with a greater variety of actions and events. For example, one approach may be to maintain a larger library of re-usable adaptations and develop a case-based adaptation mechanism. Work is also underway on other components of our architecture. We will integrate more sophisticated narrative manager and user modeling components as they become available.

We plan to explore other potential application scenarios for VE event adaptation. We feel that this approach could offer particular benefits in VE training environments, and that there may also be applications for collaborative VEs. Eventually we also hope to support user interactions beyond navigation, so that users have the opportunity not only to view events, but also to interact and influence the scenario as well.

Narrative event adaptation appears to have potential for reducing the tension between user and author control of VE narratives. While work remains on ways to achieve more flexible and effective implementations, the positive results of narrative event adaptation to date have been encouraging, and we see many possible extensions and additional application areas for this technology.

7. REFERENCES

[1] Amerson, Daniel and Kime, Shaun. Real-Time Cinematic Camera Control for Interactive Narratives. in *The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, Stanford, CA, March 2001.

[2] Bares, William H. and Lester, James C. Intelligent Multi-shot Visualization Interfaces for Dynamic 3D Worlds, *Proceedings of the 1999 International Conference on Intelligent User Interfaces*, 1999, pp.119-126

[3] Bates, J. Virtual Reality, Art, and Entertainment. *Presence: The Journal of Teleoperators and Virtual Environments*, vol. 1, no. 1, 1992, pp. 133-138.

[4] Drucker, Steven M. and Zeltzer, David, CamDroid: a System for Implementing Intelligent Camera Control, *Proceedings of*

the 1995 Symposium on Interactive 3D Graphics, 1995, pp. 139-144.

[5] Galyean, Tinsley. Narrative guidance of Interactivity. Ph.D. Thesis, School of Architecture and Planning, Massachusetts Institute of Technology. 1995.

[6] Gershon, Nahum and Page, Ward. What Storytelling Can Do For Information Visualization, *Communications of the ACM*, vol. 44, no. 8, 2002, pp. 31-37.

[7] He, Li-wei, Cohen, Michael F., Salesin, David H. The Virtual Cinematographer, *Proceedings of the 23rd annual conference on Computer Graphics and Interactive Techniques,* 1996, pp.217-224

[8] Laurel, Brenda. *Computers as Theatre.* Addison-Wessley Publishing. NY. 1991.

[9] Magerko, Brian. A Proposal for an Interactive Drama Architecture, *AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, 2002

[10] Mancini, Clara. From cinematographic to hypertext narrative, *Proceedings of the 11th ACM conference on Hypertext and Hypermedia,* 2000, pp.236-237.

[11] Marsella, Stacey C., Johnson, W. Lewis, LaBore, Catherine, Interactive Pedagogical Drama, *Proceedings of the fourth international conference on Autonomous agents*, 2000.

[12] Morie, J. F. et. al. Emotionally Evocative Environments for Training, Army Science Conference, 2002.

[13] Murray, Janet. *Hamlet on the Holodeck*. MIT Press, Cambridge, MA, 1998.

[14] Riedl, Mark, Saretto, C.J., and Young, Michael R. Managing interaction between users and agents in a multi-agent storytelling environment, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems,* 2003 (in press).

[15] Sgouros M. Nikitas, Papakonstantinou George, Tsanakas Panagiotis, Dynamic Dramatization of Multimedia Story Presentations, *Proceedings of the 1997 International Conference on Intelligent User Interfaces,* 1997, p.87-94

[16] Steiner, Karl E. Adaptive Narrative Virtual Environments, in *Computer Graphics and Multimedia: Applications, Problems, and Solutions*, ed. John DiMarco (Idea-Group Publishing, Inc.), in press.

[17] Steiner, Karl E. and Moher, Thomas G. Encouraging Task-Related Dialog in 2D and 3D Shared Narrative Workspaces, in *Proceedings ACM Conference on Collaborative Virtual Environments (CVE '02)*.

[18] Swartout, W., et al., Toward the Holodeck: Integrating Graphics, Sound, Character and Story. *Proceedings of 5th International Conference on Autonomous Agents*, Montreal, Canada, June 2001.

[19] Vardi, Guy. Navigation scheme for interactive movies with linear narrative, *Proceedings of the 10th ACM conference on Hypertext and Hypermedia,* 1999, pp.131-132.

[20] Weal, Mark J. et al. Building Narrative Structures Using Context Based Linking, *Proceedings of the 12th ACM conference on Hypertext and Hypermedia,* 2001, pp.37-38.

[21] Weyhrauch, Peter. Guiding Interactive Drama. Ph.D. Thesis, School of Computer Science, Carnegie Mellon University. Technical Report CMU-CS-97-109.

[22] Young, Michael R. An Overview of the Mimesis Architecture: Integrating Intelligent Narrative Control into an Existing Gaming Environment, *AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment*, 2001.