Steps Towards a Computational Theory of Interactive Narrative in Virtual Worlds

R. MICHAEL YOUNG

Department of Computer Science

North Carolina State University

Raleigh, NC, 27695

young@csc.ncsu.edu
919.513.3038

In this paper, we set out a basic approach to the modeling of narrative in interactive virtual worlds. This approach adopts a bipartite model taken from narrative theory, in which narrative is composed of *story* and *discourse*. In our approach, story elements – plot and character – are defined in terms of plans that drive the dynamics of a virtual environment. Discourse elements – the narrative's communicative actions – are defined in terms of discourse plans whose communicative goals include conveying the story world plan's structure. To ground the model in computational terms, we provide examples from research under way in the Liquid Narrative Group involving the design of the Mimesis system, an architecture for intelligent interactive narrative incoroprating concepts from artificial intelligence, narrative theory, cognitive psychology and computational linguistics.

Interactive narrative, artificial intelligence, planning, cognitive models, suspense, discourse generation

Introduction

The number and type of computer system using interactive 3D interfaces continue to grow as the processing power of commercial graphics cards increases. While a significant portion of the most popular virtual worlds applications are in the $9 Billion per year interactive entertainment market, it is now common for users to interact with virtual worlds in applications ranging across simulation, training, education and social interaction. Many of these environments, especially those that are focused on entertainment, exploit informal adaptations of narrative techniques drawn from conventional narrative media in their design. Even so, there is little knowledge about the nature of *computational* models of narrative that might enable the construction of computer systems that automatically create and manage engaging narrative experiences.

To increase our understanding of the nature of human interaction with narrative and virtual worlds, the Liquid Narrative Group at North Carolina State University is adapting models of narrative from narrative theory, computational linguistics and cognitive psychology, integrating these approaches with techniques from artificial intelligence in order to create interactive intelligent narrative virtual worlds systems. In this paper, we describe the approach to this project that we are currently undertaking, giving examples to some of our completed work and pointing towards the project's future goals.

A Multi-Disciplinary Approach

Our research seeks to adapt and extend existing work in artificial intelligence to account for specific story-oriented applications within 3D virtual environments. Our approach is based on methods first developed in three fields related by, in part, a common interest in narrative: cognitive psychology, computational linguistics and narrative theory. These three fields each contribute to our work along different lines. For example, cognitive psychologists provide representations of the mental models that people form when understanding narratives, and their results suggest ways that narratives interact with those models to achieve particular affective or cognitive results in those people that read, watch or interact with them. Linguists, particularly those that have studied discourse, suggest computational approaches to the telling of a narrative and how specific communicative resources can be brought to bear to convey the elements of a story in a coherent and engaging manner. Narrative theorists, while perhaps the least computational of the three, provide the basic framework for our computational models, indicating the fundamental parts of a story and it's telling, and suggesting the relationships and constraints that should hold between the data structures and algorithms intended to produce the same kinds of structures operating inside a virtual environment.

In the following sections, I will describe specific instances where approaches from these fields are being modified to create *liquid narratives*, interactive experiences within 3D virtual worlds where the action in these worlds are a) structured as a compelling story, b) generated automatically each time a user interacts with the world, c) conveyed to the user using effective narrative discourse and d) significantly altered by and adapted to the user's interaction during the course of the story. Because the structure of narrative as determined by narrative theorists plays such a central part in the design of our overall architecture, I discuss it here on its own, and the structure of story and discourse that I set out will serve as the organizing metaphor for the remained of the paper.

Story and Discourse

Narratologists have provided an extensive characterization of narrative and its elements, describing the fundamental building blocks used by an author to create a compelling story. This work, however, is analytic in nature and does not directly lend itself to a computational model capable of being used in a generative capacity. A central challenge of any computational approach that seeks to operationalize concepts from narrative theory is to determine appropriate methods to translate concepts derived from analysis into concrete, formal models capable of being put to use in the creation of an interactive virtual environment.

While a broad range of approaches to the analysis of narrative exists, our work makes use of a structure that divides a narrative into two fundamental parts -- the story and the discourse [1,2] – and we construct distinct representations and tools to manage each. From a narratological perspective, a *story* consists of a complete conceptualization of the world in which the narrative is set. This includes all the characters, locations, conditions and actions or events that take place during the story's temporal extent. Two fundamental components of a narrative – its plot and its characters – are defined within the story itself. Distinct from the story, but closely tied to it, is the narrative *discourse*. Our discourse model represents those elements responsible for the *telling* of the story, rather than containing the story elements themselves.

In our approach, the construction of a narrative discourse can itself be divided into two conceptual aspects. One aspect is the determination of both the content of the discourse and its organization. To compose a narrative discourse, an author makes choices about those elements from the story to include in the story's telling and those elements to leave out. Further, the author determines additional information about the story-world to convey to the reader (e.g., properties of relevant objects, internal properties of the story's characters). Finally, the author must organize the discourse, determining what is to be told first, what second, etc, and how the sub-parts of the discourse should arranged so as to achieve the intended communicative effects on a reader.[1]

A second aspect to the generation of narrative discourse is the selection of the specific communicative resources to be used to convey the story's elements to the reader. In a 3D virtual environment, these resources include a range of media, from voice-over narration to 3D camera control to soundtrack. The work that we describe here focuses on the generation of coherent, cinematic camera control, though our results are applicable to aspects of communicative actions across media.

---

[1] In this article, we will use the term *reader* to refer to a person that experiences a narrative when that reference is independent of the narrative's medium (e.g., when reading a novel, watching a film or interactive with a virtual world). I will use the term *user* when referring to the specific context of a person interacting with a narrative within a virtual environment.

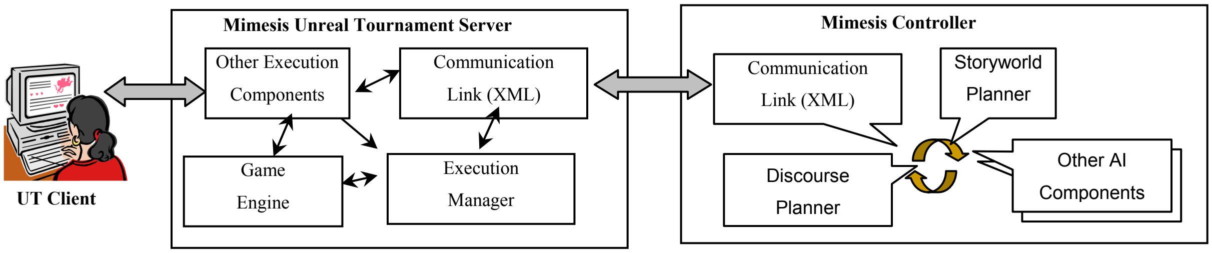

Figure 1. The Mimesis system architecture.

Creating The Narrative's Story

Change is central to the nature of narrative. In most narratives, change is initiated by the characters that populate the story world as they perform actions to achieve their individual and collective goals. We are developing a *plan-based* model of the control of characters within virtual worlds and constructed an architecture, called Mimesis [3,4], that uses this model to generate story-level plans for characters operating within a narrative. The Mimesis system integrates a suite of intelligent control tools with Unreal Tournament (UT), a commercially available 3D graphical game engine. While UT is well-suited as an engine for building conventional 3D interactive game titles, the representation of the environments that it models does not match well with those typically used by AI researchers. Like most game engines, UT's internal representation is *procedural* -- it does not utilize any formal or declarative model of the characters, setting or the actions of the stories that take place within it. Consequently, direct integration of intelligent software components is not straightforward.

To facilitate this integration, Mimesis overrides UT's default mechanisms for controlling its virtual environment, using instead a client/server architecture in which low-level control of the game environment is performed by a customized version of the game engine (called the *Mimesis Unreal Tournament Server,* or *MUTS*) and high-level reasoning about narrative structure and user interaction is performed remotely by an intelligent control element (called the *Mimesis Controller,* or *MC*). The architecture is presented in Figure 1 and described briefly below. Figure 2 shows several examples of virtual worlds built for and controlled by the Mimesis architecture.

Within Mimesis, the MC acts essentially as a story server, determining the narrative elements of the user's experience within the virtual world. The MC is responsible for both the generation of a story (in the form of a *story-world plan* characterizing all character actions that are to be performed within the environment) and the maintenance of a coherent narrative experience in the face of unanticipated user activity. At start-up, the MC first generates a story-world plan [5] and then encodes the relevant plan structure into an XML formatted message, transmitting the message to the MUTS via a socket connection. Each discrete, primitive action in the planning representations used by the MC is mirrored by a functional definition in the game engine that implements the action's semantics. Upon receiving the XML message, the *Execution Manager*, the element within MUTS responsible for driving the story's action, builds a directed acyclic graph whose nodes represent individual actions in the plans and whose arcs define temporal constraints between actions' orderings. The Execution Manager uses one-to-one mappings from the action types of the nodes in this graph to game engine functions and from the parameters of each action to instances of game engine objects in order to construct function calls that will drive the appropriate animations and state changes within the virtual world.

The narrative planner searches for a story plan – a sequence of actions to be carried out by the characters in the story (including the character controlled by the user) that will both satisfy the goals of the story and provide an engaging narrative arc. The discourse planner takes as input the story plan and a library of communicative actions that can be used by the game engine to convey the unfolding action of the story. These actions might include directives for a 3D camera controller [6], instructions for narrative voice-overs, the use of background music, etc. The discourse planner creates an action sequence containing directives to be carried out not by characters in the story world but by the game engine's interface resources and intended to be executed concurrently with the story plan itself.

The plan structures that we employ are produced by the DPOCL planner [7]. DPOCL plans are composed of *steps* corresponding to the actions that characters carry out within a story; in DPOCL, each step is defined by a set of *preconditions*, the conditions in the world that must hold immediately prior to the step's execution in order for the step to succeed, and a set of *effects*, the conditions in the world that are altered by the successful execution of the action. In addition to a set of steps, a DPOCL plan contains a set of *temporal constraints* defining a partial temporal ordering on the execution of the plan's steps and a set of *causal links* connecting pairs of steps. Two steps s_1 and s_2 are connected by a causal link with associated condition c (written $s_1 \rightarrow^{c} s_2$) just when c is an effect of s_1 and a precondition of s_2 and the establishment of c by s_1 is used in the plan to ensure that c holds at s_2.

Further, DPOCL plans contain information about the hierarchical structure of a plan, similar to the representation used by hierarchical task network (HTN) planners [8]. In DPOCL plans, steps are organized into a hierarchy. The lowest level of the hierarchy contains primitive steps, those steps directly executable by code running within the virtual world. Steps higher up in the hierarchy represent abstract actions; each abstract step in a plan is connected to a set of more primitive steps that represent a sub-plan for achieving the abstract action's effects. Because action sequences within narratives are often *episodic* that is, because they follow common patterns of action, these hierarchical structures are particularly amenable to representing story fragments.

Overall, DPOCL plans contain a fair amount of causal, hierarchical and temporal structure. This structure is added by the planning algorithm in order to ensure the plan's formal properties. Specifically, this structure is added in order to guarantee that a plan is *sound*, that is, that, when executed, the plan is guaranteed to achieve its goals. Further, this structure and the means by which it is added during planning ensures that the approach satisfies a restricted form of *completeness*, that is, DPOCL is guaranteed to find a plan for a particular problem if such a solution exists.

Adopting a plan-based model of story structure has three key advantages. First, the formal properties of the planning algorithm – specifically the planner's soundness and completeness – guarantee that the plans contain adequate structure to effectively control the story world's virtual environment. These properties of DPOCL make the plans it produces well-suited for use in controlling the execution of a virtual environment [4, 9]. [2]

A second benefit to the use of plans to drive a narrative is in its correspondence to a user's mental model of the story it defines. Recent work [10, 11] suggests that this structure as well as the techniques used by the DPOCL algorithm to create it, make for effective models of human plan reasoning. Our recent work indicates that the core elements of the plans match up with the models of narrative structure defined and validated by psychologists [12].

Despite the correspondence between plan structures, the virtual world's execution model and the user's cognitive models of a story, there are limits to the expressivity of existing plan representations when applied to narrative.   For instance, one central aspect of narrative that is not currently modeled in plan representations is that of narrative conflict. Conflict is an identifiable concept in most planning representations, though by design it does not exist in the plans that they produce. This type of conflict is associated in DPOCL with *threats* to causal links [13]. A step s_k threatens a causal link $s_i \rightarrow^c s_j$ just when $\neg c$ is an effect of s_k and s_k might occur during the temporal interval between s_i and s_j. Threats may be unintentionally created as DPOCL builds a plan, but are automatically detected and removed when they arise during the course of the planning process.

While plans containing threatened causal links might serve to model the types of conflict found in narrative, DPOCL's current mechanism for removing threats results in the production of plans that are threat-free. This is an attractive property for conventional planning systems, ensuring that the plans DPOCL produces are sound. Conflict in narrative, however, is not something to be avoided. Narrative conflict, while perhaps requiring as much effort to manage in the production of a story as threat resolution does in plan generation, must be added to a story-world plan, maintained during its production and execution and balanced with the successful achievement of character goals to create a coherent experience for the user.

---

[2] While plans that are always sound has certain advantages, soundness itself is not always appropriate for narrative plans, as I describe below.

Conflict arising from multiple agents acting to achieve disparate goals has been modeled in the context of adversarial planning. However, most approaches to adversarial planning typically create models of conflict that arise from models of a collection of agents, each acting independent from its adversaries, taking whatever action that moves it closer to the achievement of its own desires and/or that interferes with the progress of its opponents. This model of planning is well-suited for competitive games such as Bridge or for real-world military simulations, but does not match well with the nature of conflict in narrative.

To a user, narrative conflict may appear similar to conflict modeled in adversarial planning. But the nature of narrative conflict imposes a range of additional restrictions on a story-world plan that differentiates the problem of its generation from that of conflict in other planning contexts. Conflict in a story-world plan must be *balanced* between protagonist and antagonist; no one character must ever overwhelm another until the final climax of the story. Further, conflict must be *measured*, so that conflict a) occurs at key intervals rather than at every opportunity and b) supports rising action, that is, the gradient increase of tension due to conflict, with peaks and valleys of conflict intensity leading to a major conflict and its resolution as the culmination of the story. In short, narrative conflict goes beyond the conflict introduced locally by individual characters in a story; it must be directly integrated into the overall plan structure so that it contributes to the story's coherence.

To account for the features of narrative conflict we describe above, we are investigating techniques for creating plans with conflict based on the structural properties of the story-world plans themselves. For instance, we exploit the hierarchical operators within our plan representation to encode action sequences where conflict is intended to occur. Boundaries between sub-graphs in a story-world plan, created by transitions between scenes corresponding to two distinct abstract story actions, often serve as appropriate locations within the story to increase or decrease conflict. Characteristics of individual steps in plans also signal those actions that are particularly good targets for threats in a particular story. As an example, Trabasso and Sperry [14] indicate that the degree to which one action is causally connected to other acts within the same story directly affects the significance of that action in a story's mental model. In DPOCL plans, this connectedness corresponds to actions that have greater numbers of incoming causal links (that is, many actions in the story have been executed in order for this act to succeed) or to actions that have greater numbers of outgoing causal links (that is, many actions in the story depend upon this action for their successful execution).

By relating a precise computational model of action to a mental model of narrative, we are able to make predictions about the cognitive and affective consequences to a user that is experiencing the execution of specific kinds of plan structures. For instance, we have been exploring the role of plans and planning on suspense, an essential property of conventional narrative forms such as the film or novel. While suspense can be of many forms and arise in many kinds of situations, we have focused on suspense deriving from a user's knowledge of the unfolding plans of a narrative's characters. In particular, we focus on suspense arising from the anticipation of future events and their consequences on the goals of the narrative's protagonist (whether that protagonist is the user or another character within the narrative).

Recent work in cognitive psychology [15, 16, 17] has considered the role of narrative structure in the creation and maintenance of this type of suspense in film and literature. Gerrig and Bernardo [15] suggest that people who read fiction act as problem-solvers, continuously looking for solutions to the plot-based dilemmas faced by the characters in a story-world. Their work indicates that a reader's suspense is dependent upon the number of solutions that she can find to the protagonist's problem: suspense is greater when there are fewer solutions accessible.

Our approach approximates the problem solving activity that a user performs when seeking solutions to plot-related problems as planning in a space of story-world plans. Our cognitive model employs the model of planning as *refinement search* defined by Khambamphati, *et al* [18]. A refinement planning algorithm represents the space of plans that it searches using a directed graph; each node in the graph is a (possibly partial) plan. An arc from one node to the next indicates that the second node is a refinement of the first (that is, the plan associated with the second node is constructed by repairing some flaw present in the plan associated with the first node). In typical refinement search algorithms, the root node of the plan space graph is the empty plan containing just the initial state description and the list of goals that together specify the planning problem.

Nodes in the interior of the graph correspond to partial plans and terminal nodes in the graph are identified with completed plans (solutions to the planning problem) or plans that cannot be further refined due for instance, to inconsistencies within the plans that the algorithm cannot resolve.

We have successfully used this model to approximate the plan-based reasoning performed by readers when understanding instructional texts [11]. By characterizing the space of plans that a user might consider when solving problems faced by a protagonist at a given point in a plot, the model can be used to make predictions about the amount of suspense a user will experience at that point. To do this, we model the set of beliefs held by a user as she experiences an unfolding narrative. At any point in the narrative, a user will have a set of beliefs about that state of the story world, a set of beliefs about the goals of the story's protagonist and a set of beliefs about the action operators available to the characters acting within the story. These three elements are used to create a DPOCL *planning problem*, the specification used as input to DPOCL to create a plan (in this case, a plan that solves the protagonist's goals given the story world's current state as known by the user). DPOCL's refinement search algorithm creates not just a single plan that solves the protagonist's goals, but the space of possible plans given the problem specification. To the extent that DPOCL's configuration mirrors the user's planning process, this plan space approximates the space of solutions considered by the user when searching for solutions to the protagonist's current problems.

To determine the level of suspense a user may experience at a particular point in a story, we relate characteristics of this space to the psychological results described above. For instance, when there are a large number of successful solution plans at the leaves of the plan space graph, we would predict that user's experience of suspense would be lower than when there were few successful solutions to the current planning problem. Additional features of this space, for instance, the ratio of failed plans to total plans, may prove to be an element of an effective prediction of suspense. We are currently evaluating this model. As a result of early experimental results, we are considering extending the representation of plans to include additional representational features, such as the perceived probabilities of success for each of the plans, and are evaluating those new features for their role in the problem-solving process of users.

(a) (b) (c)

Figure 2. The images above show example worlds controlled by Mimesis, including (a) a virtual tour of the Monterey Bay Aquarium, (b) a futuristic bank robbery story and (c) the great hall of a medieval castle.

Results from this work also suggests means by which suspense can be increased or decreased. The features of a plan space that a user considers will differ depending on the planning problem she is attempting to solve. By controlling what facts the user believes about the story world at a given time, for instance, by generating camera sequences that explicitly convey some facts while explicitly eliding others, the system can, in effect, define a planning problem with the suspense properties appropriate for each point in the narrative.

Creating the Narrative's Discourse

A narrative system must not only create engaging story-world plans, it must use its resources to tell the story effectively. In this paper we discuss one particular strategy used in the effective creation of a narrative: building narrative discourse involves the central task of determining the content and organization of a sequence of camera shots that film the action unfolding within a story world.

In our work, we build on our previous research on the generation of natural language discourse to generate discourse plans for controlling an automated camera that is filming the unfolding action within a 3D story-world. In this approach, 3D camera shots and shot sequences are viewed as planned, intentional action whose effects are seen in the cognitive state of the user. Individual camera shots are treated as primitive communicative actions, multi-shot sequences and cinematographic idioms are characterized using hierarchical plan operators, and, as in conventional discourse planning, plan structure that specifies the communicative content of a discourse is created to achieve particular effects upon the mental state of the user.

Conventional discourse planners take as input a set of propositions intended to be conveyed to the user of a system, along with a model of the user's existing knowledge of the domain of discourse and a library of plan operators describing both the primitive communication actions available to the planner (e.g., typically speech acts such as INFORM or REQUEST) and definitions for a set of abstract actions and their sub-plan schemas, sometimes referred to as *recipes*. Abstract operators often specify rhetorical structure [19,20] in a discourse (e.g., when one part of a discourse stands as evidence for the claim set forth in a second part of a discourse) and their sub-plan schemas specify how more primitive collections of communicative actions can be combined to achieve the abstract act's communicative effects.

There are several important ways that this problem – and our proposed solution to it – differ from the task of discourse generation in conventional contexts. In our approach, the propositional content that the narrative discourse planner receives as input refers not just to relations that hold in the domain of discourse, but also to propositions describing the structure of the story-world plan. The task of the discourse planner is, in part, to generate camera action sequences that convey the execution of story-world plan actions to the user.

Beyond the difference in propositional content of the discourse, generating plans for 3D narrative discourse involves a range of issues. We identify three key challenges here:

- **Generating Cinematic Discourse Structure**. Cinematic discourse contains both rhetorical structure, aimed at conveying propositions about the story world to a user, but also idiomatic structure mirroring the use of patterns for shot composition used in film [21,22]. Our plan operators must represent both these aspects of discourse structure, and combine them effectively to tell the story.

- **Representing Intermediate Discourse Goals**. Conventional discourse planning views the discourse problem as one in which a plan for communicative action is built to achieve a set of goals. These goals, typically describing what the user comes to know or believe, must hold at the *end* of the discourse. In discourse plan operators, concern is given to achieving goals and (in some systems) maintaining those goals, but there are no requirements placed on the epistemic state of the user in the discourse plan's *interior*. In contrast, the experience of a narrative requires very specific epistemic states, not just at the end of the story, but throughout its execution. Suspense, for instance, may rely on the fact that the user knows key aspects of the story world not known to other characters, and that the user know these aspects at key points while the story plays out. Similarly, interesting narrative structure may rely on a user being ignorant of key facts in the world (for instance, another character's actions leading to conflict) until a precipitating event results in their revelation. Our plan operators must provide a means by which system developers can specify intermediate epistemic goals.

- **Integrating story-world and discourse-level plans.** The actions in discourse plans for narrative in virtual worlds, unlike actions in plans for textual narrative, must themselves execute. Camera actions for panning, tracking, fading, etc, all require time to play out, a physical location from which the camera films, physical objects that must be included or excluded from the field of view, etc. A particularly complicating aspect of this is that these camera actions must execute in the same the same temporal and spatial environment as the objects of the story that they must convey to the user. A knowledge representation for narrative discourse must take this shared evironment into account or risk creating suboptimal (or even inconsistent) plans.

  For instance, consider the case where camera action C_1 is responsible for filming action s_1 and camera action C_2 is responsible for filming action s_2. If s_1 completes its execution prior to s_2 beginning, then C_1 must complete its execution prior to C_2. In plans where the successful execution of C_1 depends upon the user knowing some property of the domain first established in C_2, the inconsistency must be detected and remedied. If s_1 and s_2 happen concurrently in the same spatial location, then C_1 and C_2 could be replaced by a single appropriately placed shot C_3. Without considering their colocation/co-occurrence, a planner would not be able to generate this option. Futher, if s_1 and s_2 occur in sequence in locations that are spatially adjacent, then a single shot C_4 that *tracks* (that is, moves the camera) from the location of s_1 to the location of s_2, could be used in place of a C_1 ; C_2 sequence. Again, the representation of the context for execution of both discourse and story level plans is required in order to even consider the option of such a narrative.

The Challenges of Interactivity in Story and Discourse

Control and Coherence

Our work described above takes an idealized stance in which the user is not accounted for except as a passive observer. While this assumption is entirely valid for conventional narrative media such as film or literature, the assumption is almost always invalid for narrative-oriented virtual worlds. A key feature of these worlds is the level of interactivity that they offer the user. The ability to step into the narrative world and play a character in the story, to take substantive action within the unfolding story, is a key distinguishing feature of virtual worlds and stories.

A central issue in the development of effective and engaging interactive narrative environments is the balance between coherence and control. The understandability of any narrative is determined, in part, by it's *coherence*, that is, by the user's ability to comprehend the relationships between the events in the story, both within the story world (e.g., the causal or temporal relations between actions) and in the story's telling (e.g., the selection of camera sequences used to convey the action to the user). Dramatists often refer to narrative as having a premise or point; stories are told for a reason and much of our comprehension of a story involves the construction of cognitive models that predict or explain these relationships. Systems that construct actions for telling a story should respect the story's coherence by clearly linking each action in the story world to its overall structure.

The degree of engagement by a user within an interactive narrative lies, to a great extent, with the user's perceived degree of *control* over her character as she operates within the environment. The greater the user's sense of control over her character, the greater will be her sense of presence [23], that is, the sense that she is a part of the story world and free to pursue her own goals and desires. Unfortunately, control and coherence are often in direct conflict in an interactive narrative system. To present a coherent narrative, the actions within an interactive narrative's story are carefully structured (either at design time by human designers, in the case of conventional computer games, or at run time by intelligent systems like the one we're developing) so that actions at one point in the story lead clearly to state changes necessitated by actions occurring at subsequent points in the story. When users exercise a high degree of control within the environment, it is likely that their actions will change the state of the world in ways that may interfere with the causal dependencies between actions as intended within a storyline.

Conventional forms of narrative (e.g. film and novel) resolve the issue of coherence versus control by completely eliminating control; the audience is a passive observer. Computer game developers, in contrast to film makers, introduce interactivity in their systems, but carefully limit the control exercised by the user by designing the environment so that the user's choices for action at any point reduce to a small set of options moving the user through a pre-defined branching structure. In the remainder of this paper, we discuss a technique used in the Mimesis architecture called narrative mediation which allows a degree of control and coherence that lies between that of computer games and conventional narrative media.

Managing Interactivity at the Story Level

As described above, Mimesis drives the action within its story world based on the structure of a plan produced by a narrative planner. As users issue commands for their characters to perform actions within the story world, these actions have the potential to undo conditions in the world that are critical to the success of actions in the narrative plan that have not yet executed. Consequently, before carrying out directives from the user, the corresponding actions must be checked against the narrative plan to determine how they fit with the plan's structure. Within Mimesis, each action α performed by the user is automatically characterized in one of three ways with respect to the unexecuted portion of the plan. One possibility is that α is *constituent* to the plan – α matches an action prescribed by the narrative plan for execution by the user, in which case the user is doing exactly the action that the system desires her to do in order to perform that portion of the storyline.

The second possibility is that α is *consistent* with the plan – α is not constituent and none of the effects of α interact with any of the plan's remaining structure. For example, it may be consistent if the user rotates her character in a circle in order to orient herself spatially before walking out of a room, as long as her act of walking out of the room is part of the narrative and is successfully performed during the appropriate timeframe. The third possibility is that α is *exceptional* – α is not constituent and one or more of α's effects threaten the conditions in the world required by future agent actions. Specifically, an exception occurs whenever a user attempts to perform some action α, where some effect $\neg e$ of α threatens to undo some causal link $s1$ e $s2$ between two steps, $s1$ and $s2$, with condition e, where $s1$ has occurred prior to α and $s2$ has yet to occur.

If a user performs an exceptional action, the effects of the exception on the virtual world undoes the condition of at least one causal link in the plan, invalidating some or all of the plan's subsequent structure. It is the responsibility of the system to detect exceptions when they arise and to respond accordingly in a manner that balances the need to preserve the coherence of the narrative with the need to preserve the user's sense of control. Within the Mimesis system, response to exceptions occurs in one two ways. Either the system allows the exception to occur and restructures the narrative plan mid-story, or it prevents the exception from actually executing, in effect coercing the user into compliance with the existing plan structure. We refer to this process of exception detection and response as *narrative mediation*.

In narrative mediation, exceptions are dealt with in one of two ways. The most straightforward is via *intervention*. Typically, the success or failure of an action within a virtual environment is determined by software that approximates the physical rules of the underlying story world (e.g., setting a nuclear reactor's control dial to a particular setting may cause the reactor to overload). However, when a user's action would violate one of the story plan's constraints, Mimesis can intervene, causing the action to fail to execute.

In the reactor example, this might be achieved by surreptitiously substituting an alternate set of action effects for execution, one in which the "natural" outcome is consistent with the existing plan's constraints. A control dial momentarily jamming, for instance, will preserve the apparent consistency of the user's interaction while also maintaining safe energy levels in the story world's reactor system.

The second response to an exception is to adjust the narrative structure of the plan to accommodate the new activity of the user. The resolution of the conflict caused by the exception may involve only minor restructuring of the narrative, for instance, selecting a different but compatible location for an event when the user takes an unexpected turn down a new path. Accommodation may involve more substantive changes to the story plan, however, and these types of modifications can be computationally expensive. For instance, should a user instigate a fight with a character that is intended to be a key ally later in the story or unintentionally destroy a device required to rescue a narrative's central character, considerable re-planning will be required on the part of the MC's narrative planner.

In order to handle exceptions in an interactive narrative system, the narrative planner analyzes its plans prior to execution, looking for points where enabled user actions (that is, actions whose preconditions for execution are satisfied at that point in the plan) can threaten its plan structure. When potential exceptions are identified, the planner weighs the computational cost of re-planning required by accommodation against the potential cost incurred when intervention breaks the user's sense of agency in the virtual world. The reader is encouraged to see [2] for a discussion of the technical approach used for narrative mediation within Mimesis.

Managing Interactivity at the Discourse Level

An effective narrative experience is determined by at least two factors: the successful execution of the story's actions and the effective communication about those actions to the user. When the user controls a camera that acts as her viewport into the story, the same issues regarding the balance of control and coherence come to play at the discourse level. In most 3D virtual environments, the camera is associated with the point of view of the user's character; as she moves her character's position and direction of gaze, her camera translates and rotates accordingly, changing her field of view into world and altering the set of story actions that she can observe within it.

Changes made by the user to her camera orientation will typically conflict with the structure of a narrative discourse in which the orientation of the camera is specified for at each moment of the story. Consequently, an account of user camera control that characterizes each change as either constituent, consistent or exceptional with respect to the narrative's discourse plan seems impractal. Our current work has not addressed exceptions that arise due to a user's camera control actions; we have, instead, built narrative structures that have either been non-interactive (e.g, fly-throughs) or ones that have interleaved segments of unrestricted interactivity (in which no camera control is imposed) with ones where camera control is removed from the user entirely (e.g., cut scenes).

We are approaching this problem by extending the model of a user's camera control actions and the model of the ways that those actions may contribute to the goals of a narrative discourse. This extension will take into account the changes in knowledge state caused by a user's choice of shot composition; by integrating this model into a discourse-level mediation scheme, we expect that techniques similar to the accomodation approach described in the previous section will be applicable. However, a discourse-level equivalent to intervention seems more problematic. Because camera control is tightly coupled with a user's input actions (e.g., mouse manipulation), an approach that substitutes alternate camera control actions without disrupting a user's sense of agency in the virtual world is a challenge to define.

Conclusions

In this paper, we have set out a basic approach to the modeling of narrative in interactive virtual worlds. This approach adopts a bipartite model of narrative as story and discourse in which story elements – plot and character – are defined in terms of plans that drive the dynamics of a virtual environment. Narrative discourse elements – the narrative's communicative actions – are defined in terms of discourse plans whose communicative goals include conveying the story world plan's structure. To ground the model in computational terms, we have provided examples from the Mimesis system, including features of the system that have already been implemented and features that are under development, targetted at the theoretical division we have set out above. While there are many possible means to approach a story-and-discourse model of interactive narrative, our goal is to demonstrate the effectiveness of this model using the Mimesis system as a testbed; our initial results, mentioned here and in the work we cite, are encouraging.

Acknowledgements

We are grateful to the many students that have contributed to the Liquid Narrative Group and the development of the Mimesis system. Significant contributors to the work described in this paper include Dan Amerson, Mark Branly, D'Arcey Carol, Arnav Jhala, Shaun Kime, Milind Kulkarni, R.J. Martin, Mark Riedl, C.J. Saretto and Sirish Somanchi. This work was supported by National Science Foundation CAREER award #0092586.

References

[1] Chatman, S. Story and Discourse: Narrative Structure in Fiction and Film, Cornell University Press 1990.

[2] Emmot, K. Narrative comprehension: A Discourse Perspective, Oxford, 1999.

[3] Young, R. M. And Riedl, M. Towards an Architecture for Intelligent Control of Narrative in Interactive Virtual Worlds. In: Proceedings of the International Conference on Intelligent User Interfaces, January, 2003.

[4] Riedl, M., Saretto, C.J. and Young, R. M. Managing interaction between users and agents in a multi-agent storytelling environment In: Proceedings of the Second International Conference on Autonomous Agents and Multi-Agent Systems, June, 2003.

[5] Young, R. M. Notes on the Use of Plan Structures in the Creation of Interactive Plot. In: *The Working Notes of the AAAI Fall Symposium on Narrative Intelligence*, Cape Cod, MA, 1999.

[6] William H. Bares , James C. Lester, Intelligent multi-shot visualization interfaces for dynamic 3D worlds, In: Proceedings of the 4th international conference on Intelligent user interfaces, p.119-126, January 05-08, 1999, Los Angeles, California, United States

[7] Young, R. Michael, Martha E. Pollack and Johanna D. Moore, Decomposition and causality in partial-order planning, in Proceedings of the Second International Conference on AI and Planning Systems, Chicago, IL, pages 188-193, July, 1994.

[8] Sacerdoti, E. A Structure for plans and behavior. Elsevier, New York, 1977.

[9] Cavazza, M., Charles, F. & Mead, S., Planning characters' behaviour in interactive storytelling. In Journal of Visualization and Computer Animation 2002; 13(2): 121-131.

[10] Rattermann, M. J., Spector, L., Grafman, J., Levin, H. and Harward, H. Partial and total-order planning: evidence from normal and prefrontally damaged populations, Cognitive Science 2002; 25(6); 941-975.

[11] Young, R.M. Cooperative Plan Identification: Constructing Concise and Effective Plan Descriptions. Proceedings of the National Conference of the American Association for Artificial Intelligence, pages 597-604. Orlando, FL, 1999.

[12] Christian, D. and Young, R. M. Comparing Cognitive and Computational Models of Narrative Structure. Liquid Narrative Technical Report TR03-001, Liquid Narrative Group, Department of Computer Science, North Carolina State University, Raleigh, NC. 2003.

[13] Penberthy, J. and Weld, D. UCPOP: A sound, complete, partial order planner for ADL. In Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning, 103--114, October 1992.

[14] T. Trabasso and L. Sperry. Causal relatedness and importance of story events. Journal of Memory and Language, 1985, 24:595-611.

[15] Gerrig, R., and Bernardo, D. Readers as problem-solvers in the experience of suspense. Poetics 1994; 22; 459-472.

[16] Ohler, P. and Niedling, G. Cognitive modeling of suspense-inducing structures in narrative films. In Suspense: Conceptualizations, theoretical analyses and empirical explorations. Vorderer, et al, eds, Lawrence Erlbaum. 1996.

[17] Comisky, P., and Bryant, J. Factors involved in generating suspense. Human Communication Research 1982; 9(1):49-58.

[18] Khambamphati S, Knoblock, C and Yang, Q., Planning as Refinement Search: A Unified Framework for Evaluating Design Tradeoffs in Partial-Order Planning in Artificial Intelligence vol. 76(1-2), 167-238, 1995.

[19] Mann and Thompson. Rhetorical structure theory: A theory of text organization. In Livia Polyani, editor, The Structure of Discourse. Ablex Publishing Corporation, 1987.

[20] Moore, J. and Paris, C. Planning text for advisory dialogues: Capturing intentional and rhetorical structure, Computational Linguistics 1993; 19(4): 651-695.

[21] Arijon, Daniel. Grammar of the film language, Communication Arts Books, Hastings House: New York, 1976

[22] Monaco, J. How to Read a Film: The Art, Technology, Language, History and Theory of Film and Media, Oxford University Press, 1981.

[23] Lombard, M. & Ditton, T. At the heart of it all: the concept of presence. Journal of Computer-Mediated Communication 1997; 3(2).