A Proposal for an Interactive Drama Architecture

Brian Magerko

University of Michigan
1101 Beal Ave.
Ann Arbor, MI 48109
magerko@umich.edu
www.soargames.org

Abstract

We approach creating an interactive drama from the viewpoint that having a human author involved in the creation of the drama is a valuable component. Given that there is this author, we suddenly find opposing forces inherent to the system: the amount of interaction we desire to give the User of the system vs. the amount of dramatic control we wish to give the Writer. In this paper, we propose IDA, an Interactive Drama architecture, which is based on the principle of having a real-time Director agent intelligently guiding the User's experience towards the Writer's dramatic goals, while at the same time allowing as much User flexibility as possible.

Introduction

Traditional storytelling can be an enriching experience for the creator as well as the reader; everyone loves a good story. The writer of a story has the opportunity to tell some tale that they desire to put into form; they have the gift of expression. The reader of the story is on the receiving end of this communiqué. They have an enriching experience comprised of an interesting storyline, characters, and settings. We are proposing to contribute to an alternative approach to telling stories: interactive drama. Brenda Laurel's definition has come to aptly characterize this yet-to-be-realized art form (Laurel 1986):

> An "interactive drama," then, is a first-person experience within a fantasy world, in which the User may create, enact, and observe a character whose choices and actions affect the course of events just as they might in a play. The structure of the system proposed in the study utilizes a playwriting expert system that enables first-person participation of the User in the development of the story or plot, and orchestrates system- controlled events and characters so as to move the action forward in a dramatically interesting way.

This definition brings up several important points, the most important being that in an interactive drama, *the User is the character*. We want to tailor a dramatic experience specific to the User interacting with it. There are two main visible approaches to altering the story being told. The first approach would be to create believable, autonomous characters in a story setting. The User would be placed in an environment with these richly-defined characters, and the story would unfold differently based on how the User interacted with the other characters in the story (Sgorous 1999). The plot would evolve out of local changes in the state of the world made by the characters. The second approach would involve an omniscient, God-like director that would be able to control the plot globally, tuning and tweaking the state of the world according to the User's interactions with the system (Weyhrauch 1997; Mateas & Stern 2000).

An alternative to this black and white view is a synthesis of the two approaches; the director could make observations on what should change, and give the characters some higher-level goal or change in behavior that they would then carry out to fulfill the director's change in plot. Our group is currently focusing on this approach. This paper will present a blueprint that we have constructed for interactive drama that we call *IDA* (Interactive Drama Architecture).

This project has both technical as well as artistic goals in mind. The technical objective that we hope to accomplish is to *heuristically guide a User's interactions within a plot that has been abstractly defined by the Writer of the system.* By formulating a suitable heuristic approach to plot direction, we can have a better understanding of how our technology compares to past built systems in interactive fiction (Sgorous 1999; Weyhrauch 1997) as well as current approaches (Young 2000; Szilas 2001; Mateas & Stern 2000). We want to understand how to use the information available to advance (or direct) plot. Artistically, we desire *to offer the User a rich, yet flexible story-intensive*

experience that will be different, yet still dramatically interesting, with different interactions. We also desire to allow the storywriter (referred to hereafter as the *Writer*) enough flexibility to *tell the story they desire to tell*, while at the same time presenting the User with an experience that gives them a maximal amount of control over what happens in the story. This balance of Writer flexibility vs. User flexibility has arisen as a major problem in the design of interactive drama systems in the recent past. Trying to find the balance between these two goals is a higher-level goal of this project (see Figure 1**).**

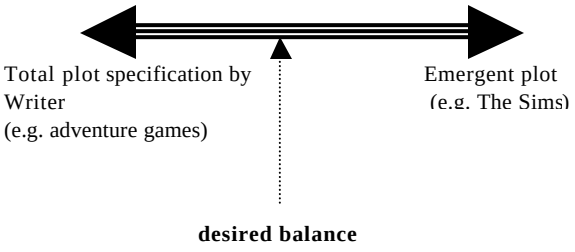

Figure 1: The desired balance of plot control in our

Design

As with any new art form or technology, the exact parameters we work with can be vague and differ from artist to artist, or from engineer to engineer. In order to make it clear how we intend to view this problem, we propose several requirements for an interactive drama (ID):

- **Variability**: As the User behaves differently each time he uses the system, how do we allow the story to be different in some significant way?
- **Writer flexibility**: Does the system allow the Writer to tell the story they want to tell? Does it allow for a full range of dramatic expression?
- **Balance**: Does the system hit a good balance between Writer flexibility and User flexibility?
- **Transparency**: How do we encourage the User to follow a particular destiny without having him feel forced into it?

We propose this architecture, IDA, as a means to meeting the challenges of these requirements. IDA has several key components: the environment, the character, the User, the Writer, and the Director. The relationships between these components and the information shared between them are detailed below in Figure 2. We also offer a description of the key elements that we have been focusing on, namely the environment, Writer, and Director.

The Environment

We desire the ability to tell a story in a richly defined, 3D environment, likening the experience to literally giving the User the control and viewpoint over the protagonist in a movie. Our current work is on utilizing the Unreal Tournament engine to give us a test bed for work in plot direction as well as synthetic characters, though this paper will be mainly focusing on plot direction and representation. Work is currently be done in using this technology to build Haunt II, a haunted-house mystery story that includes the User as the protagonist (Laird, et al 2002).

Plot Representation

We have committed ourselves to the notion that the Writer wishes to directly communicate some vision in the story that is told in an interactive drama system. We have attempted to give the Writer a reasonable amount of control in plot specification, while keeping under consideration that we also desire to give the User a significant amount of control over how the drama unfolds. Our specification language is a first-order logic with several components for each scene. We view a scene as the smallest dramatic unit the Writer uses to move the story forward as a whole (i.e. a complete story is a higher-level unit than an act, which in turn is higher-level than a scene, which is above the atomic unit, the beat) (Field 1994).

- Initial state
- Required events
- Background knowledge
- Content constraints
- Temporal constraints

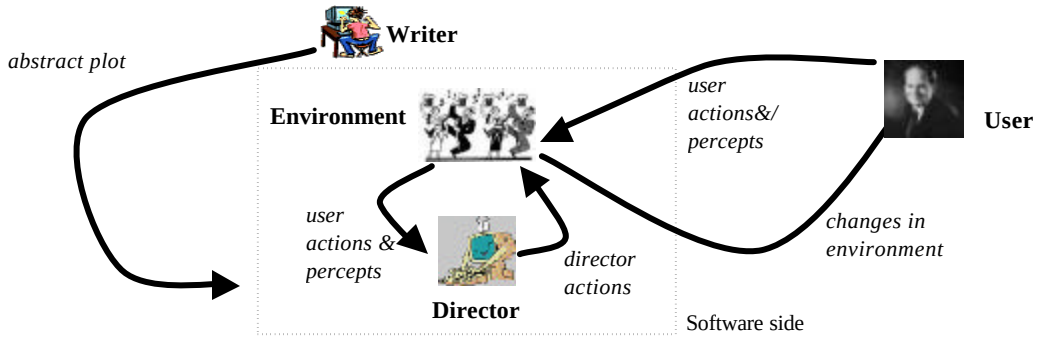

Figure 2: The Interactive Drama Architecture

The *initial state* for a scene is the setup for the scene. It details where the characters are, what they have on their person, what they are doing, etc. This can be likened to a richer, more detailed representation of the typical setup for a scene in a screenplay. A screenplay may contain some blocking, the general mood of the scene, even perhaps the initial action that is taking place as the scene begins. All of these things can be contained within our initial state. We also make an assumption that to some degree our characters are richly-defined, directable agents with goal-based behavior. Given this assumption, our representation may also feasibly contain modifications to the agent's state, such as the objects it carries with it, knowledge it has, the agent's goals, etc.

The *required events* are a conjunction of goal conditions for the scene. They are statements about the state of the world that must be true for the current scene to end and for the next scene to begin. Rounding out the definition of a scene is *background knowledge*, a collection of operators with preconditions and effects. These events contain variables that may be locally or globally defined. The global definition of a variable across scenes allows us to ensure that the plot content will be consistent. For example, if the Writer constructs a plot that requires the person who gets closest to the User to wind up betraying him in the end, then we want to be sure that whichever variables should represent that person are all bound to the same character. Therefore, we see the need for allowing global variables, in both the required events as well as in all other components, to be shared between scenes. These events and operators can be viewed as a formalization of the dramatic concept of a *beat*, the atomic event in a dramatic scene that actually creates some change in the world.

We may also constrain the partial ordering of these required events with *temporal constraints*. Temporal constraints on goal events indicate events that must be fulfilled before others, just as a POP-style planning language may allow (Kautz and Selman 1992).

The *content constraints* limit the binding of the variables that are used in the required events. For the moment, we are not allowing disjunctive constraints for simplicity in the representation. However, we will consider them further down the road as we explore the problem further.

Variability

The ideal interactive drama would have a richly-defined world, with an infinite (or at least sizably large) amount of actions available to the User (just as in the real world). No matter what actions the User executed, an interesting dramatic experience would develop over time. Each time that the User started the system anew, different interactions would lead to pleasantly varied dramatic experiences. We view variability to mean that there are many possible complete stories that may be told in an entire experience with an ID.

Putting aside possible variance due to Director actions, our design philosophy focuses on separating interactivity into two types of variability: temporal variability and content variability. We define *temporal variability* as allowing time to be the key variable for the flexibility in an interactive experience. Who? What? Where? and Why? are all static aspects of the story; the only variable is when scenes actually occur. Pete Weyhrauch's PhD thesis (Weyhrauch 1997) at Carnegie Mellon is a prime example of introducing temporal variability into a first-person dramatic experience to produce an interactive drama. We

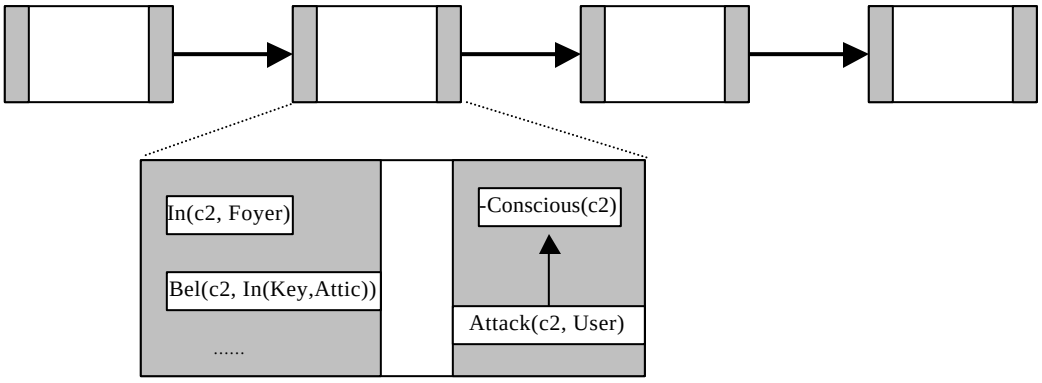

Figure 3: An example of scene ordering. The close-up is a view of how the scene is defined in terms of an initial state and required events. The other elements discussed, temporal and content constraints, as well as background knowledge, are not depicted. They are inputted into the Director without any processing (such as inputting them into a planning system to construct a plan).

define *content variability* as allowing scenes to be defined in an abstract fashion, giving flexibility into what specifically happens in that scene. Who the murderer is where a certain scene takes place….who falls in love with the antagonist….who aids the protagonist at the end….all of these things are examples of abstract concepts in a plot that deal with the content of the story.

Note that we've decided to address the problem of content variation before concerning ourselves with temporal variability. Scene ordering is specified by the Writer and is static in this representation. In order to make progress in either of these types of variability, we will focus solely on one of them and leave the other as a control variable in our experimentation.

Plot Direction

Now that we have shown how to loosely specify a plot, how do we take advantage of this sudden flexibility? In other words, what happens when it's unclear how a variable should be instantiated when it needs to be? Since we have a firm notion of how we wish to represent plot, our representation naturally leads us to this fundamental question. Answering it will be a primary goal of our future experimentation. For now, we have a working concept of how we can approach it:

- Annotate each possible object, character, etc. with attributes (whatever they may be) and then heuristically choose the binding for each variable.
- Have a default binding for each variable. This would put more work on the designer, but at the same time would allow the Designer to input an instantiation of the story as the default.
- Randomly choose a binding.

In this order, we will take these steps as our first hypothesis to experiment with for plot direction in our representation. We define plot direction to mean having the current world state S, some dissimilar desired world state S', and executing some action in the world to indirectly encourage the User to enter S'. Note that this definition has two distinct parts: choosing a desired state, and then executing some action to encourage it.

The act of choosing a desired state at any point is a fairly straightforward procedure, save for the heuristic design. The Director agent must recognize that the User has achieved the last desired state and there is a new event that it must consider. The Director must then choose which event should happen next (since we have allowed a loose temporal ordering of the events within a scene), and with what variable bindings. We propose that the Writer design a heuristic to make this decision, much like ID systems before (Weyhrauch 1997; Sgorous 1999).

Once the Director desires a state S, it must then execute some action to encourage the User to enter that state. We have yet to design an intelligent method of generating these actions and therefore have put this task into the hands of the Writer as part of the plot specification. As mentioned before, our representation language has obvious parallels to a planning language. We specify an initial state, operators with preconditions and effects (background knowledge), and goals (required events). We specify a partial ordering of the operators and required events with our temporal constraints. Given a particular scene, one can imagine using a POP planning system (like Graph-plan or SAT-plan) to create a feasible partial-order plan (or a set of plans) given our constructs. A resulting plan would represent the space of possibilities for what the User could do in the world and how the world would change. Given this representation, we can now create a relationship between User behavior and direction by annotating our plan with directions that will encourage the various operators to be executed.

We can also see the Director taking on more tasks than what we have described above, which we may call variable monitoring. Another function of the Director may be knowledge monitoring, which would involve the following task: if the User or Director is involved in some event in the world that affects other synthetic character's knowledge bases, then the appropriate knowledge should be shared. Any information that is not directly obtained from characters' percepts is relevant. A special case of this role, goal monitoring, would be the same as knowledge monitoring, except with making sure the characters' current goals are consistent with the plot.

Mateas and Stern's architecture proposal is similar to ours in many respects (Mateas & Stern 2000). They focus on an in-scene plot director that manages the characters' behavior on a beat-by-beat basis. Within a scene, characters are passed human-assembled joint plans describing the coordinated activity required of all the characters for a particular beat, with a specified sequence of beats representing a scene. Both approaches can be viewed as a hierarchically decomposed representation. In their case, they decompose the characters' behavior in the description of how a beat should unfold and how their actions can coordinate with others'. In our work, we specify how a scene should unfold in more high-level, abstract terms, passing off decomposition of a given Director action (such as a character decomposing a new goal or action given by the Director) into the hands of the given synthetic character.

By giving the synthetic characters a stronger role in determining their own behavior (i.e. specifying their lower-level behavior internally as opposed to externally in the plot representation), our approach offers several benefits. We have more flexibility in the kinds of directions we can

consider giving a character. For example, if a scene is running on too long for some reason, the Director could possibly give the synthetic characters a goal to "hurry up and end the scene," (much like real actors often do in improvised scenes). Having autonomous characters means that their behavior is modular; we do not have to rewrite total behaviors from scene to scene or even story to story. We can also rely on the characters not to 'drop dead' if some error occurs in the system, or some unforeseen state is entered and the Director is not sure what to do at a particular moment. While there are arguments for having weaker, less autonomous characters, there is certainly enough of an argument for our proposed approach to warrant exploration.

Discussion

The end result of using this representation is a skeletal plot specification created by the human Writer. One of the strengths of this approach is that it guarantees that any possible dramatic device can be employed within the story; the Writer can directly add anything he desires in the plot. Take the plot of the film The Sixth Sense (my apologies to those who have not seen this film and might take this discussion as a plot spoiler). The plot veritably takes on two different forms. One plot line is a story about Bruce Willis' character helping out a young, troubled boy. The other plot line doesn't take shape until the very end of the film: the audience learns, quite by surprise, that all along Bruce Willis has actually been a ghost the entire film, trying to complete his unfinished work in the world. An ID based on emergent plot would likely have a hard time developing a plot such as the one exhibited in The Sixth Sense. Even if one were created, it would undoubtedly mean hand-coding this specific dramatic device into the system, which an approach that is being explored in other systems (Bringsjord & Ferrucci 2000; Sgorous 1999). Encoding all of the elements, special rules, etc. of drama seems to be an insurmountable task. Putting the creation of dramatic content into the hands of the Writer frees us from this worry.

Given that we can vary the content of the plot, what determines which states should be encouraged / discouraged? How does the User's actions affect this decision? Though we are far from a complete answer to these questions, our specification language gives us a good starting point to answer them. Once our architecture is fully built, we will be able to begin experimentation on the effectiveness of this representation and how well it answers the questions above. It should be noted that working our way up to a completely built architecture will be a time and labor-intensive task. However, there are several comments to be made in general about our work thus far and what questions we have still to address:

The types of dramatic devices that may be used are limited only by the Writer's imagination. By leaving the intelligence of plot specification up to the Writer, we're guaranteeing a much broader and deeper coverage of dramatic possibilities. Implementing a system that created an emergent plot based on recognizing and encouraging interesting dramatic possibilities is a different approach to building an ID. The introduction of drama is done through encoding dramatic principles into a set of rules. While this may be a more autonomous approach, it is a much more intensive task to encode all of the necessary elements of drama into a general ID rather than handing some of the plot writing over to a human Writer, which is our approach.

It is unclear what kinds of content variability are possible, both for our system and in general. Given the nature of our representation, we can be assured that the "Who?" part of a scene may be variable for instance, as we've shown in our examples. What about deeper features of the plot though? Is this representation sufficient enough to model variability in complex relationships between characters? in complicated sequences of actions? Or are these aspects of the plot perhaps too complicated to allow much flexibility? We believe that once we have built a functional, working system, we will be able to vigorously experiment with both our representation and our approach to direction without compromising our architectural design.

Not only are there states the Director will wish to encourage, but there may be states to discourage as well. As in computer games or real life, there may be some desires that cannot be upheld if certain actions are executed in the world. For instance, if I desire to run for president of the United States, but then commit and get convicted of a felony, then it is impossible for me to achieve my goal of the presidency. There are irreversible actions, actions the User may be able to execute that are irreversible and damaging to the plot, which we may wish the Director to acknowledge and help avoid. How does the Director achieve this?

Using a user model may be an important input to the Director. Given that drama is heavily involved in the goal-based behavior of the characters involved (e.g. the protagonist and antagonist may likely have conflicting goals), it stands to reason that having a guess about the User's current goals in a situation would be beneficial to directing the plot. There has been already work done in a similar domain called a Multi-User Dungeon (or MUD), where a dynamic Bayes network was built to reason about the User's current goals in the dungeon (Albrecht, et al 1997). Reasoning about how the User's goals and knowledge relate to the goals and knowledge of other characters would help us have a more in-depth reasoning about how the plot should unfold. As Syd Field wrote, "Drama is conflict," (Field 1994). Without knowing the goals of all the characters in the plot (including the User as

a character), then we can't adequately introduce conflict into our plot based on their goals. This brings us to our last discussion point, intelligent characters.

What types of Director actions should be possible? Up until this point, we have mainly focused on Director actions changing some aspect of the states of a given synthetic character. However, one can imagine a situation where an auditory or visual cue in the environment may yield the desired effect of encouraging or discouraging the User to enter a particular state. While we are not focusing on these possible types of direction for now, they certainly are possible ways of affecting the User and should be addressed in the future.

What types of experimentation should we conduct? Simply put, how do we in fact know when we're done? What makes an interactive drama more successful than other approaches to it? than other iterations of the same approach? To fulfill our artistic goals, we can only assume that user enjoyment has to play a key part in assessment of the overall success of the project. The variables we may consider in experimenting are the many possible technical alterations we can make to the system's design: how much autonomy is given to the characters, how the representation language is specified, what kind of direction the Director could enact, etc. After constructing a fully-functional prototype, we will be approaching these issues and evaluating how they affect User enjoyment, as well as system load, communication bandwidth between agents, and other technical concerns. Designing a methodology for user testing will be an important step in showing that a system that fits our definition of an interactive drama is actually a "successful" example of one.

How do directable, intelligent characters fit into this framework? Given that there are intelligent agents running around the environment, each with their own beliefs, goals, and knowledge, what sort of information needs to be communicated between the agents and the Director? How does this affect the kinds of direction the Director is capable of? These questions will be approached as both our research on plot direction as well as As-Sanie's work on directable agents progresses (As-Sanie 2002).

References

Albrecht, D. W., Zukerman, I., Nicholson, A. E. and Bud, A. 1997. Towards a Bayesian Model for Keyhole Plan Recognition in Large Domains. User-Modeling: Proceedings of the Sixth International Conference. 367-376. Sardinia, Italy. : CISM.

As-Sanie, M. 2002. Directable Synthetic Characters. Artificial Intelligence and Interactive Entertainment: Papers from the 2002 AAAI Spring Symposium. Menlo Park, Calif.: AAAI Press.

Bringsjord, S. & Ferucci, D. A. eds. 2000. Artificial Intelligence and Literary Creativity. Mahwah, New Jersey: Lawrence Erlbaum Assoc., Inc.

Field, S. eds. 1994. Screenplay: The Foundations of Screenwriting. New York: Del Trade.

Kautz, H. and Selman, B. 1992. Planning as Satisfiability, Proceedings of the 10th European Conference on Artificial Intelligence. 1-12. Vienna, Austria.

Laird, J., As-Sanie, M., Bachelor, B., Benninghoff, N., Enam, S., Stokes, D., Jones, B., Kerfoot, A., Lauver, C., Magerko, B., Sheiman, J, Stokes, D., Wallace, S. 2002. A Test Bed for Developing Intelligent Synthetic Characters. Artificial Intelligence and Interactive Entertainment: Papers from the 2002 AAAI Spring Symposium. Menlo Park, Calif.: AAAI Press.

Laurel, B. 1986. Toward the Design of a Computer-Based Interactive Fantasy System. PhD diss., Drama Department, Ohio State University.

Mateas, M. and Stern, A. 2000. Towards Integrating Plot and Character for Interactive Drama. In *Working Notes of the Social Intelligent Agents: The Human in the Loop Symposium.* AAAI Fall Symposium Series. Menlo Park, CA: AAAI Press.

Sgorous, N. M. 1999. Dynamic generation, management, and resolution of interactive plots. Artificial Intelligence 107: 29-62.

Szilas, Nicolas 2001. A New Approach to Interactive Drama: From Intelligent Characters to an Intelligent Virtual Narrator. Artificial Intelligence and Interactive Entertainment: Papers from the 2001 AAAI Spring Symposium. 72-76. Menlo Park, Calif.: AAAI Press.

Weyhrauch, P. 1997. Guiding Interactive Drama. PhD diss., Computer Department, Carnegie Mellon University.

Young, R. Michael 2000. Creating Interactive Narrative Structures: The Potential for AI Approaches. The Working Notes of the AAAI Spring Symposium on Artificial Intelligence and Interactive Entertainment. Menlo Park, Calif.: AAAI Press.