Improving Narrative Consistency in Planning-Based Interactive Storytelling

Leandro Motta Barros
Universidade do Vale do Rio dos Sinos
Programa Interdisciplinar de
Pós-Graduação em Computação Aplicada
São Leopoldo, RS, Brazil
leandromb@unisinos.br

Soraia Raupp Musse
Pontifícia Universidade Católica do Rio Grande do Sul
Programa de Pós-Graduação em
Ciência da Computação
Porto Alegre, RS, Brazil
soraia.musse@pucrs.br

Abstract

Interactive Storytelling (IS) systems are an emerging class of interactive entertainment applications with emphasis in narrative aspects. This paper deals with stories that have plots centered on a riddle, and describes a mechanism that can be introduced to planning-based IS in order to make the generated stories follow a specified tension arc. Results from experiments with users suggest that the proposed method is capable to bring improvements to these systems.

Introduction and Related Work

Games for computers and consoles are established as the leading form of interactive digital entertainment. Technologically, the development of the field of computer games is impressive. A more conceptual analysis, though, shows that most games developed in the past decades have explored fundamentally the same concepts: hand-eye coordination, puzzle-solving and resource management (Crawford 2004). Countering this reality, several researchers have been exploring aspects normally treated as marginal by the majority of today's games, like the dramatic and narrative aspects.

The field that aims to create interactive applications capable to generate consistent narratives is called Interactive Storytelling (IS). Researchers on this field are exploring different approaches in order to achieve the goal of seamless conciliation of interactivity and narrative consistency. One of these approaches is based on the use of planning algorithms to generate the sequence of events that compose the story. Planning-based IS has characteristics that help to satisfy the interactivity requirements of IS systems, but has some problems concerning the narrative consistency. In this paper, we propose a mechanism that intends to improve these systems by making the generated stories follow a tension arc defined by the storyworld author. This mechanism is focused on stories that follow the Riddle master plot (Tobias 1993).

Planning has been used to develop narrative-related systems since the seminal works on non-interactive story generation, like UNIVERSE (Lebowitz 1985). In this work, Lebowitz automatically generates plot outlines that resemble soap operas through the use of planning techniques and libraries of plot fragments and characters.

The present work is based on the use of STRIPS-like planning systems, an approach introduced by Charles *et al.* (2003). In a previous work (Barros & Musse 2005), we built on these ideas, focusing on ways to make the generated stories more narratively consistent.

The design of the interactive drama Façade, by Mateas and Stern (2003), shows a lot of care in aspects related to generating highly consistent narratives. In particular, they try to ensure that the generated stories follow a certain tension arc. Our paper also presents a mechanism that aims to make the generated stories follow an author-defined tension arc, but our work (which is in the context of planning-based IS and centered on riddles) is considerably different.

Architecture Overview

Before delving into the method we propose for respecting tension arcs, we present a brief overview of our IS architecture (details can be found in (Barros & Musse 2005)). We see the story as a sequence of actions performed by characters that can transform the world in a way imagined by the storyworld author. The player controls the story protagonist. Actions performed by the player may render the current plan invalid. Whenever this happens, the planning algorithm is used to create a new plan, thus adapting the story to the player's actions. Our current implementation uses the Metric FF planning system (Hoffmann 2003).

World state, represented as a set of predicates, is changed whenever characters perform actions. This includes both actions performed by NPCs (Non-Player Characters) and by the player-controlled story protagonist. Our prototype displays the story to the player with animated characters in a 3D environment (Figure 2). The player executes actions with the mouse: clicking in an object (or character) brings a list of actions related to it.

Riddles and Tension Arcs

The level of tension in a story in function of time results in a curve called *tension arc*. The goal of our work is to introduce a mechanism that makes the generated stories follow an author-defined tension arc. We are dealing with stories that follow the Riddle master plot (Tobias 1993), *i.e.*, stories in which there is a mystery to be solved, like "whodunits". Unfortunately, Tobias is very brief when discussing tension or tension arcs for this kind of story. Actually, he only mentions that "the tension of your riddle should come from the conflict between what happens as opposed to what seems to have happened."

Based on this Tobias' assertion, we modeled tension arcs for riddle stories starting from the premise that the tension raises as new clues are discovered or revealed, such that the knowledge that the player has about the events (what seems to have happened) approaches to the truth (what actually happened). The culmination of the tension is the moment in which the player needs just the final clue to solve the riddle. The story ending coincides with the instant in which the player's knowledge reaches what actually happened, that is, it is the moment in which the riddle is solved. Thus, in the proposed model, the player knowledge about the riddle and the tension are equivalent.

Formalizing these ideas, the tension arc of a story is a function of time, designated by $K(t)$. For instance, if the player starts the story without any clue that may help to solve the riddle, the player's knowledge (and therefore the tension) would be zero, *i.e.*, $K(0) = 0$. The value of K is incremented whenever a clue is revealed to the player. The magnitude of this increment depends on the importance of the clue for the solution of the riddle. Generically, we say that for a clue c, K will be incremented by k_c. The values of k for each clue are to be chosen by the storyworld author. The desired tension arc, designated by $K^*(t)$, is defined by the storyworld author. We can also define a function, $D(t) = K^*(t) - K(t)$, that measures the discrepancy between these two curves. Figure 1 illustrates these concepts. Our current implementation requires $K^*(t)$ to be composed of increments at discrete instants, hence the "stairs-like" appearance of the desired tension arcs seen of figures 1 and 3.

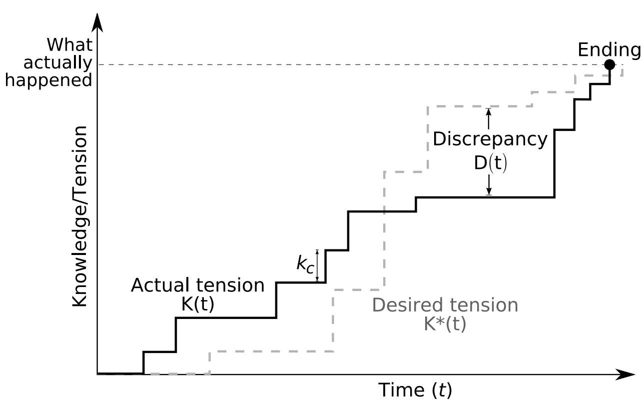

Figure 1: Model of tension arc.

Ideally, we would like to have $D(t) = 0$ for all possible values of t. In practice, however, there always will be some difference between the two arcs. So, to quantify the "total discrepancy" observed in a story, we introduce another quantity,

$$E = \int_{t=0}^{T} |D(t)|\,,$$

where T is the total story duration. This way, the intuitive notion of respecting a tension arc can be formally defined as minimizing E.

In the proposed model, the value of $K(t)$ just changes when a clue is revealed. Hence, to reduce the discrepancy $D(t)$ (and consequently minimize E), it is necessary to increase or decrease the rate at which clues are revealed. To achieve this goal, we dynamically change the level of participation that NPCs have in the discovering of clues. In other words, if the player is facing difficulties to find clues ($K(t) < K^*(t)$), the other story characters should help. Likewise, if the player is finding clues too quickly ($K(t) > K^*(t)$), the NPCs should work less to help. We do this by varying the cost of actions performed by NPCs (denoted c_{NPC}), as described below. The cost of the protagonist actions is constant: $c_{\text{prot}} = 100$.

We consider that, for any instant t, there exists a value, $c^*_{\text{NPC}}(t)$, that represents the ideal cost for the NPC actions at that time. This value is considered the ideal in the sense that using it will incur the generation of a plan in which the NPCs will act in the right amount to bring the story back to its correct tension level. The amount of NPC participation is dependent on the current discrepancy $D(t)$, so $c^*_{\text{NPC}}(t)$ is defined as a function of $D(t)$.

In order to properly define $c^*_{\text{NPC}}(t)$, we conducted several experiments with the Metric FF planning algorithm and Ugh's Story 2 (described in "Results"). We generated and analyzed several plans using various values for c_{NPC}. Based on this, we identified some good values for NPC action costs given some values of $D(t)$. These points empirically found could be well fitted by the function $c^*_{\text{NPC}}(t) = \lfloor -1.875D(t) + 75 \rfloor$, which was used in the experiments reported in this paper.

While running, every 15 seconds, our system compares the value of the NPC action cost used to create the current plan to the value of $c^*_{\text{NPC}}(t)$. If the difference between them is over 25%, or if the absolute difference is above 20, a new plan is generated using the current value of $c^*_{\text{NPC}}(t)$ for the NPC actions cost.

Results

The ideas presented in this paper were implemented in our prototype and tested in practice with a storyworld named "Ugh's Story 2". Everything happens in Neandertown, a small village inhabited by cavemen. Neandertowners worship a statue that used to lay on an altar near the caves in which they live. One day, Alelugh, the priest, was found unconscious near the altar, hit by a club, and the statue was stolen. The player plays the role of Ugh, a caveman hired to discover the author of such crime. Besides Ugh and Alelugh, two characters take part in the story: Egonk and Brokung. As the story unfolds, the player discovers that Egonk is very egotist and selfish (and could have stolen to have the statue just for himself) and Brokung is bankrupt (so, he could have stolen the statue to sell it). As more clues are revealed, the player eventually finds out that Brokung stolen the statue and sold it to Salesorg, the traveling salescaveman. An image of the storyworld is shown in Figure 2.

We let five users play Ugh's Story 2. Three of them used the system as described in this paper; the other two used

Figure 2: Our prototype running Ugh's Story 2.

a modified version of the prototype that didn't include the method for respecting the desired tension arc. For the first group, the resulting total discrepancies (E) were 9337.5, 4434.9 and 5972.4 (mean: 6581.6). For the second group, values were 63804.0 and 31251.8 (mean: 47527.9). While the sample is too small to provide statistical guarantees, the reduction in the total discrepancy is clear.

Figure 3 depicts the run of *Ugh's Story 2* by one of the players. Here, we can see a practical example of the system acting to bring the tension back to its desired level. From 3:08 to 6:00, the player advanced slower than expected, making the actual tension level considerably lower than the desired. As result, a new plan is generated, using a lower value for c_{NPC}. The lower c_{NPC} resulted in the presence of more actions performed by NPCs in the new plan. Consequently, starting from time 6:00, the NPCs start to work harder to help finding clues. This culminates with a NPC telling a clue he has found to the player, at 7:00. This brings the tension level to a value very close to the desired.

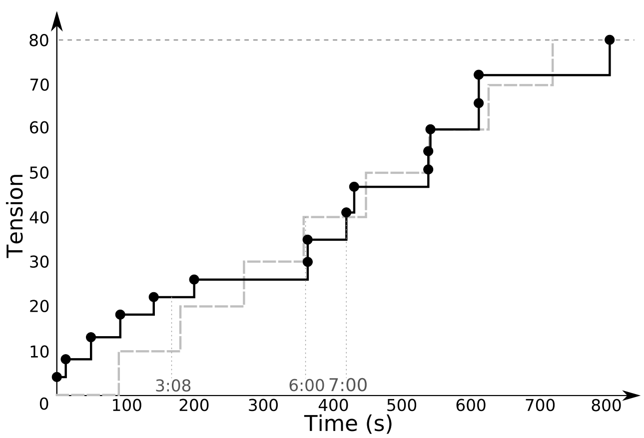

Figure 3: Tension arc for the story discussed in "Results" (continuous black line). The desired tension arc is shown for reference (dashed gray line). Dots indicate moments in which the player discovered clues.

Final Remarks

The field of Interactive Storytelling is still incipient and deals with a very complex topic. Several different approaches are being used by different research groups. This work is based on the use of a planning algorithm to define the sequence of events that compose the story.

Planning algorithms are typically created in the context of Artificial Intelligence (AI) research, which has different goals than those in IS. AI is mostly concerned with "hard" and precise goals like optimality. On the other hand, the narrative goals in IS are more subtle, not easily defined formally. Hence, nothing guarantees that plans created by these algorithms correspond to good stories.

In this paper, we presented a technique that can be used in the context of planning-based IS to confront this fact. When the generated stories closely follow an ideal tension arc, the narrative will progress in the pace imagined by its author. Our results, obtained from experiments with users, suggest that the proposed method can help to advance in this direction.

There are certainly many aspects of this work that can be improved in the future. Among these possibilities, we would like to mention the creation of a more complex model of tension arcs for riddle stories. We are particularly interested in dealing with lies told by NPCs, because, when an NPC is belied, the player's knowledge about the riddle will shrink in some sense, and this is not supported by our current model.

Acknowledgments

We would like to thank Marcelo H. Borba for the cavemen models. This work is supported by the National Council for Scientific and Technological Development (CNPq – Brazil).

References

Barros, L. M., and Musse, S. R. 2005. Introducing narrative principles into planning-based interactive storytelling. In *Proc. of ACM SIGCHI Intl. Conference on Advances in Computer Entertainment Technology (ACE 2005)*, 35–42.

Charles, F.; Lozano, M.; Mead, S. J.; Bisquerra, A. F.; and Cavazza, M. 2003. Planning formalisms and authoring in interactive storytelling. In *Proc. of TIDSE'03: Technologies for Interactive Digital Storytelling and Entertainment*.

Crawford, C. 2004. *Chris Crawford on Interactive Storytelling*. Indianapolis, USA: New Riders Games.

Hoffmann, J. 2003. The Metric-FF planning system: Translating "ignoring delete lists" to numeric state variables. *Journal of AI Research, special issue on the 3rd International Planning Competition* 20:291–341.

Lebowitz, M. 1985. Story-telling as planning and learning. *Poetics* 14:483–502.

Mateas, M., and Stern, A. 2003. Façade: An experiment in building a fully-realized interactive drama. In *Game Developers Conference, Game Design track*.

Tobias, R. B. 1993. *20 Master Plots: and how to build them*. Cincinnati, USA: Writer's Digest Books.